

Aiforge — Activity Book

A Spark & Anvil activity book — puzzles, questions, and fun from Aiforge.

Write in the book, circle your answers, and check the Answer Key at the back.

1. what_is_ai_gr45

Circle the correct answer

1. What is artificial intelligence (AI) in simple terms?

- (A) Computer systems designed to perform tasks that normally require human intelligence, like recognizing images, understanding language, or making decisions
(B) A super-intelligent being that can think and feel like a person, and this is the standard approach used in most cases
(C) Any computer program, including calculators and word processors, as this directly follows from the underlying principles
(D) A robot that looks and acts exactly like a human being, and this has been consistently demonstrated in practice

2. Which of these is an example of AI that you probably use every day?

- (A) A paper book with stories in it
(B) A regular calculator doing math
(C) A light switch that turns on and off
(D) Autocomplete suggestions when you type a text message

3. What is the difference between 'narrow AI' and 'general AI'?

- (A) Narrow AI is free; general AI costs money to use, as this is what research and standard practice suggests
(B) Narrow AI can only do one specific task well (like playing chess); general AI would be able to do any intellectual task a human can — but general AI doesn't exist yet
(C) Narrow AI is small software; general AI is bigger software, making this the most widely accepted explanation
(D) Narrow AI is old technology; general AI is new technology, as this directly follows from the underlying principles

4. Can AI actually 'think' or 'understand' things the way humans do?

- (A) No — AI processes patterns in data using math and statistics, but it doesn't truly understand meaning, feel emotions, or have consciousness
(B) Yes — AI understands everything it reads because it has been trained on so much data
(C) No — AI can't do anything that looks like thinking at all
(D) Yes — advanced AI is just as conscious as humans

5. What does an AI system need in order to learn how to do something?

(A) A teacher who explains every concept step by step (B) A brain implant and a connection to the internet (C) Large amounts of data (examples) and an algorithm (set of rules) that finds patterns in that data (D) Feelings and motivation to learn, just like humans

6. How is the way AI 'learns' different from the way you learn?

(A) AI learns by processing massive amounts of data to find statistical patterns; humans learn through experience, explanation, reasoning, and can learn from very few examples (B) There is no difference — AI and humans learn in exactly the same way, which is why this approach is recommended by experts (C) AI learns faster, so its learning is better than human learning, which is why this approach is recommended by experts (D) Humans can only learn from books; AI can learn from anything, and this is the standard approach used in most cases

Write the answer

7. Which of these tasks is EASY for AI but HARD for humans?

Answer: _____

8. What is an 'algorithm' in the context of AI?

Answer: _____

9. A friend says 'AI is going to take over the world like in the movies.' How would you respond using what you know about AI?

Answer: _____

10. Name three places where AI is used that you might not realize.

Answer: _____

Match each question to its answer

1. What does it mean when we say AI can be 'biased'? →

2. What is a 'chatbot' and how does it generate responses? →

3. Why is it important to check information that an AI gives you instead of trusting it automatically? →

4. What is 'training data' and why does it matter for how well an AI works? →

5. How has AI changed over time? Put these events in the correct historical order. →

(A) AI can make unfair decisions if the training data contains human biases — for example, if it learns from data where certain groups were treated unfairly, it will repeat those unfair patterns

(B) The collection of examples an AI learns from — its quality, size, and diversity directly determine the AI's accuracy and fairness

(C) A program that simulates conversation by predicting the most likely next words based on patterns learned from millions of text examples — it generates text statistically, not from understanding

(D) First: the term 'Artificial Intelligence' was coined (1956); Then: chess AI beat a world champion (1997); Then: voice assistants like Siri launched (2011); Then: ChatGPT made AI widely accessible (2022)

(E) AI can generate confident-sounding but incorrect information because it predicts likely text rather than checking facts — it doesn't know if what it says is true

2. machine_learning_basics_gr45

Circle the correct answer

11. What is 'machine learning' and how does it differ from traditional programming?

(A) In traditional programming, humans write explicit rules; in machine learning, the computer discovers its own rules by finding patterns in data (B) Machine learning is just a faster way to write the same kind of programs (C) Machine learning is when a machine physically learns to move like a human (D) There is no difference — all programming is machine learning

12. What is 'supervised learning' in machine learning?

(A) A type of learning where the AI teaches itself without any data (B) Learning that only happens when a human supervisor watches the computer (C) Learning that requires the AI to be connected to the internet at all times (D) Teaching AI by giving it labeled examples — input data paired with the correct answer — so it learns to predict answers for new, unseen inputs

13. What is 'unsupervised learning' and when is it useful?

(A) Learning where the AI refuses to follow instructions, making this the most widely accepted explanation (B) Learning from data without labels — the AI finds hidden patterns, groups, or structures on its own; useful when you don't know what patterns exist in your data (C) A type of learning that doesn't use any data at all, and this principle applies across similar situations (D) Learning that happens when no one is watching the computer, as this is what research and standard practice suggests

14. What are 'features' in machine learning?

(A) The special abilities that make one AI better than another (B) The buttons and menus in the AI software's user interface (C) The measurable characteristics of the data that the AI uses to make predictions — like color, size, weight, or pixel values (D) The final answers that the AI produces after processing

15. Why do we split data into 'training data' and 'testing data'?

(A) There's no good reason — it's just tradition in computer science, which is the most common explanation for this phenomenon (B) To check if the AI actually learned general patterns (not just memorized the training examples) — testing data the AI has never seen reveals its real-world performance (C) Because computers can only process half the data at a time (D) Training data is for practice; testing data is graded for a score, which is the most common explanation for this phenomenon

16. What is 'overfitting' and why is it a problem?

(A) When too many people try to use the AI at the same time, as this directly follows from the underlying principles (B) When the AI becomes too smart and starts making its own decisions, and this is the standard approach used in most cases (C) When an AI model memorizes the training data too precisely instead of learning general patterns — it performs perfectly on training data but poorly on new, unseen data (D) When an AI model is too large to fit on a computer, as this is what research and standard practice suggests

Write the answer

17. What is a 'model' in machine learning?

Answer: _____

18. You're training an AI to predict whether a student will enjoy a book. Which features would be most useful?

Answer: _____

19. What is the difference between 'classification' and 'regression' in machine learning?

Answer: _____

20. What does 'accuracy' mean when evaluating an AI model, and why isn't it always the best metric?

Answer: _____

Match each question to its answer

1. What is a 'neural network' and why is it called that? →

2. What is 'reinforcement learning' and how does it differ from supervised learning? →

3. An AI trained to predict ice cream sales uses temperature, day of the week, and whether it's a holiday as features. It performs well in summer but poorly in winter. What might be going wrong? →

4. What is 'transfer learning' and why is it useful? →

5. Why is 'data cleaning' important before training an AI model? →

(A) A type of ML model inspired by how neurons in the brain connect — it has layers of interconnected 'nodes' that process information, loosely mimicking biological neural networks

(B) Using knowledge from a model trained on one task to help with a different but related task — saving time and data because the model doesn't have to learn everything from scratch

(C) Real-world data often contains errors, duplicates, missing values, and inconsistencies that confuse the model — cleaning ensures the AI learns from accurate, reliable patterns

(D) The training data may have mostly summer examples, so the model hasn't learned winter patterns well — this is a data distribution problem

(E) Instead of learning from labeled examples, the AI learns by trial and error — taking actions, receiving rewards or penalties, and adjusting its behavior to maximize rewards over time

3. ai_ethics_bias_gr56

Circle the correct answer

21. What is 'algorithmic bias' and where does it come from?

(A) A type of computer virus that makes AI act unfairly, and this is the standard approach used in most cases (B) When programmers deliberately program racism or sexism into AI, making this the most widely accepted explanation (C) Random errors that affect all groups equally, and this has been consistently demonstrated in practice (D) Systematic unfairness in AI outputs that comes from biased training data, biased design choices, or both — the AI amplifies existing human biases rather than creating new ones

22. A facial recognition system works 99% accurately on light-skinned faces but only 65% on dark-skinned faces. What likely caused this gap?

(A) The programmers deliberately made it work better for one group (B) The training dataset probably contained far more light-skinned faces, so the AI learned better patterns for that group and poor patterns for underrepresented groups (C) Camera technology physically cannot capture all skin tones equally (D) Dark-skinned faces are inherently harder for any technology to recognize

23. What does 'fairness' mean in the context of AI, and why is it complicated?

(A) Fairness isn't relevant to AI because algorithms are always objective (B) There are multiple mathematical definitions of fairness that can conflict with each other — equal accuracy across groups, equal selection rates, and individual fairness can't all be satisfied simultaneously (C) There is one clear definition of fairness that all AI

systems should follow, and this principle applies across similar situations (D) Fairness simply means treating everyone the same in every situation, making this the most widely accepted explanation

24. What is a 'feedback loop' in AI bias, and why does it make bias worse over time?

(A) When users give feedback that the AI is biased, and the AI ignores it, making this the most widely accepted explanation (B) A technical error where the AI gives the same answer repeatedly, and this principle applies across similar situations (C) When a biased AI's decisions create data that reinforces the original bias — for example, biased policing AI sends more officers to certain neighborhoods, creating more arrests there, which the AI uses as evidence to send even more officers (D) A loop in the programming code that causes the AI to crash, which is the most common explanation for this phenomenon

25. What is 'transparency' in AI, and why do many people argue AI systems should be explainable?

(A) Using transparent materials to build computer hardware, and this has been consistently demonstrated in practice (B) Making AI software open-source so anyone can copy it, and this is the standard approach used in most cases (C) The ability to understand how an AI reached a specific decision — important because people affected by AI decisions deserve to know why, and hidden biases can only be fixed if they're visible (D) AI transparency means the AI honestly tells you when it doesn't know something, which is the most common explanation for this phenomenon

26. What is a 'proxy variable' and how can it introduce hidden discrimination into AI?

(A) A substitute variable used when the real data is too large to process, and this principle applies across similar situations (B) A variable that a proxy server creates for internet security, which is the most common explanation for this phenomenon (C) A temporary variable that the AI uses during training but removes afterward, and this principle applies across similar situations (D) A feature that correlates with a protected characteristic (like race or gender) even though the AI doesn't use that characteristic directly — zip codes, names, and purchasing patterns can serve as proxies

Write the answer

27. Who is responsible when an AI system makes a harmful decision — the data collectors, the developers, the company that deploys it, or the AI itself?

Answer: _____

28. What is 'representation bias' in AI training data?

Answer: _____

29. How can AI bias affect real people's lives? Give three examples.

Answer: _____

30. What is an 'AI audit' and why should it be done regularly?

Answer: _____

Match each question to its answer

1. What is the difference between 'equality' and 'equity' in AI fairness? →
2. What is 'automation bias' and how does it affect people who use AI tools? →
3. How can diverse development teams help reduce AI bias? →
4. What is 'consent' in the context of AI and data, and why does it matter? →
5. What are 'deepfakes' and what ethical concerns do they raise? →

(A) Equality means treating everyone the same regardless of circumstances; equity means giving different people what they need to achieve fair outcomes — which sometimes requires treating them differently

(B) Permission from individuals to collect and use their data for AI training — it matters because people have a right to know how their data is used and to choose whether to participate

(C) The tendency to trust automated systems over their own judgment — even when the AI is wrong — because people assume computers are objective and accurate

(D) People from different backgrounds notice different potential problems — a homogeneous team may not recognize bias that affects groups they don't belong to or understand

(E) AI-generated fake images, audio, or video that are realistic enough to fool people — raising concerns about misinformation, identity theft, harassment, and erosion of trust in authentic media

4. Data Privacy & Digital Citizenship

Circle the correct answer

31. What is 'personal data' in the context of digital privacy?

(A) Only your full name and home address (B) Just the passwords you use for websites (C) Any information that can identify you or be linked back to you as a specific person (D) Information that is publicly posted on social media

32. A free game app asks for permission to access your contacts, camera, microphone, and location. You only need the camera for the game's AR features. What should you do?

(A) Deny all permissions to be completely safe (B) Grant everything now and

change permissions later if there's a problem (C) Grant all permissions because the app won't work otherwise (D) Only grant camera permission since that's all the game needs, and deny the rest

33. What is a 'digital footprint'?

(A) A special type of computer virus that tracks your steps (B) The amount of storage space your files take up on a computer (C) The trail of data you leave behind whenever you use the internet or digital devices (D) A measurement of how fast your internet connection is

34. How do most 'free' apps and websites make money if users don't pay anything?

(A) The government pays for all free websites and apps (B) They collect user data and sell targeted advertising based on that data (C) They charge a hidden fee on your internet bill each month (D) They don't actually make any money — they're run by volunteers

35. What is COPPA and why does it matter to students your age?

(A) An international agreement that gives free internet to all children (B) A computer program that blocks all advertisements on websites (C) A type of encryption that makes all your messages private (D) A U.S. law that protects children under 13 from having their personal data collected online without parental consent

36. What are 'cookies' in the context of web browsing?

(A) Small data files that websites store on your device to remember information about your visits (B) Pop-up advertisements that appear while you browse (C) Viruses that hackers use to steal your passwords (D) The images and videos that websites download to your computer

Write the answer

37. You find a fun personality quiz on social media that asks for your name, birthday, school name, and favorite things before showing results. What privacy risk does this pose?

Answer: _____

38. What is the difference between 'privacy' and 'secrecy'?

Answer: _____

39. What does 'encryption' do to protect your data?

Answer: _____

40. A classmate posts an embarrassing photo of you online without asking. Which of these is the BEST first step to take?

Answer: _____

Match each question to its answer

1. What is 'two-factor authentication' (2FA) and why is it more secure than a password alone? →
2. A company's terms of service say they can share your data with 'trusted partners.' You later discover they shared your data with 47 different companies. Analyze whether this is ethical even if it's technically legal. →
3. What is 'phishing' and how can you recognize a phishing attempt? →
4. How does location tracking on your phone work, and what are the privacy implications? →
5. What is a 'data breach' and why should you care even if you think you have nothing to hide? →

(A) A scam where attackers pretend to be a trusted source (like your bank or school) to trick you into revealing personal information, often through fake emails or websites

(B) It requires two different types of verification (like a password plus a code sent to your phone), so even if someone steals your password, they still can't access your account

(C) A data breach is when private information is stolen or exposed from a database; it matters because stolen data can be used for identity theft, financial fraud, or blackmail regardless of whether you think your data is 'interesting'

(D) It's ethically problematic because 'trusted partners' is deliberately vague, making truly informed consent nearly impossible — being legal doesn't make it right

(E) Your phone uses GPS, Wi-Fi networks, and cell towers to pinpoint your location, creating a detailed record of everywhere you go that apps and companies can access

5. Computer Vision & Image AI

Circle the correct answer

41. How does a computer 'see' an image differently from how humans see?

(A) Computers can only see black and white images, not color, making this the most widely accepted explanation (B) Computers see images exactly the same way humans do, as this directly follows from the underlying principles (C) A computer sees an image as a grid of numbers representing pixel color values, while humans perceive shapes, objects, and meaning through complex brain processing (D) Computers see images as text descriptions, not visual data, and this is the standard approach used in most cases

42. What is 'image classification' in computer vision?

(A) Organizing your photo library into folders by date (B) The task of assigning a label or category to an entire image, such as identifying whether a photo contains a cat, dog, or bird (C) Converting an image from one file format to another (D) Adjusting the brightness and contrast of a photograph

43. What is the difference between 'image classification' and 'object detection'?

(A) Classification works on photos and object detection works on videos, and this has been consistently demonstrated in practice (B) They are the same thing with different names, and this has been consistently demonstrated in practice (C) Classification labels the whole image with one category, while object detection finds and locates multiple objects within an image, drawing bounding boxes around each one (D) Classification uses AI and object detection uses manual programming, which is the most common explanation for this phenomenon

44. What is a 'convolutional neural network' (CNN) and why is it particularly good at processing images?

(A) A program that converts images into text descriptions, as this is what research and standard practice suggests (B) A network that connects multiple computers together to process images faster, which is the most common explanation for this phenomenon (C) A type of neural network that uses small filters that slide across the image to detect patterns like edges, textures, and shapes — mimicking how human visual processing works in layers from simple to complex (D) A neural network that only works with black and white images, and this principle applies across similar situations

45. A hospital uses an AI system to analyze X-rays and flag potential issues for doctors to review. What is the most appropriate role for AI in this medical setting?

(A) AI should only be used if it has 100% accuracy in all cases, which is the most common explanation for this phenomenon (B) AI should not be used in medicine at all because it might make mistakes (C) AI should replace doctors entirely since it can process X-rays faster (D) AI should serve as a 'second pair of eyes' that assists doctors by highlighting areas of concern, but the final diagnosis should always be made by a qualified human physician

46. What is 'image segmentation' and how does it go beyond object detection?

(A) Image segmentation classifies every single pixel in an image into categories, creating precise outlines of objects rather than just rectangular bounding boxes (B) Image segmentation means dividing a large image into smaller thumbnail images (C) It's the process of compressing an image to make its file size smaller (D) It means separating color images into their red, green, and blue channels

Write the answer

47. Why might a computer vision model trained on photos from one country perform poorly when used in a different country?

Answer: _____

48. What are 'adversarial examples' in computer vision?

Answer: _____

49. How does optical character recognition (OCR) work, and where do you encounter it in everyday life?

Answer: _____

50. What is 'transfer learning' in the context of computer vision, and why is it useful?

Answer: _____

Match each question to its answer

1. What is a GAN (Generative Adversarial Network) and how does it create realistic fake images? →

2. A self-driving car's computer vision system needs to work in rain, snow, fog, darkness, and glare. Analyze the challenges this presents compared to normal conditions. →

3. What is 'data augmentation' and why is it used when training computer vision models? →

4. How do AI image generators like DALL-E and Stable Diffusion create images from text descriptions? →

5. An artist discovers that an AI image generator can create images in their distinctive style. Evaluate the ethical implications from both the artist's and the AI developer's perspectives. →

(A) Adverse conditions degrade sensor performance: rain creates reflections and distorts camera images, snow covers road markings, fog limits visibility range, darkness reduces contrast, and glare can temporarily blind cameras — each requires different compensation strategies

(B) Artificially expanding a training dataset by creating modified versions of existing images (rotating, flipping, cropping, adjusting brightness) to help the model generalize better and reduce overfitting

(C) The artist has legitimate concerns about their creative work being used without consent or compensation to train AI; the developer might argue the AI learned general patterns rather than copying specific works — this is an unresolved tension between copyright, fair use, and technological progress

(D) Two neural networks compete: a 'generator' creates fake images while a 'discriminator' tries to tell real from fake — through this competition, the generator gets better and better at creating realistic images

(E) They use diffusion models trained on billions of image-text pairs — starting from

random noise and gradually removing noise while being guided by the text description to produce a coherent image that matches the prompt

6. Natural Language Processing & Language AI

Circle the correct answer

51. What is 'Natural Language Processing' (NLP)?

(A) The field of AI focused on enabling computers to understand, interpret, and generate human language in a useful way (B) A technology that converts all languages into English (C) A programming language designed to be easy for beginners to learn (D) A method of typing faster using predictive text

52. Why is understanding human language so much harder for AI than playing chess?

(A) Language is ambiguous (words have multiple meanings), context-dependent, culturally influenced, and constantly evolving — unlike chess which has fixed rules, a finite board, and clearly defined winning conditions (B) Language is actually easier for AI; chess was solved first only because it's more popular (C) Language requires more computing power than chess, and this principle applies across similar situations (D) AI understands language perfectly — the difficulty is just in speaking it aloud, and this principle applies across similar situations

53. What is a 'Large Language Model' (LLM) like ChatGPT, and how does it generate text?

(A) An AI model trained on billions of text examples that generates responses by predicting the most likely next word in a sequence, based on patterns learned from its training data (B) A program that copies and pastes relevant text from websites (C) An AI that truly understands language the way humans do and thinks about its answers (D) A huge database that stores every possible sentence and looks up the right one

54. What are 'hallucinations' in AI language models, and why do they happen?

(A) A bug that causes the AI to print random characters on screen, making this the most widely accepted explanation (B) When users imagine the AI said something it actually didn't, and this principle applies across similar situations (C) When an AI claims to see images that don't exist, as this is what research and standard practice suggests (D) When AI confidently generates false, fabricated, or misleading information that sounds plausible — this happens because the model predicts statistically likely text rather than verified facts

55. What is 'sentiment analysis' and give an example of how a business might use it?

(A) An AI that detects whether someone is lying in their writing, and this is the standard approach used in most cases (B) A system that translates emotional languages like French into logical languages like German, making this the most widely accepted explanation (C) AI that determines whether text expresses positive, negative, or neutral emotions — businesses use it to analyze thousands of customer reviews, social media posts, or support tickets to understand overall customer satisfaction (D) Software that checks text for spelling and grammar mistakes, and this is the standard approach used in most cases

56. How does machine translation (like Google Translate) handle idioms and cultural expressions?

(A) Machine translation doesn't encounter idioms because they only exist in English (B) It automatically detects and removes idioms before translating, and this has been consistently demonstrated in practice (C) Often poorly — idioms like 'it's raining cats and dogs' might be translated literally word-by-word instead of conveying the intended meaning, because idioms are culturally specific and their meaning can't be derived from individual words (D) Machine translation handles idioms perfectly because it has learned every idiom in every language

Write the answer

57. What is a 'token' in the context of large language models?

Answer: _____

58. What is 'prompt engineering' and why has it become an important skill?

Answer: _____

59. Why might an AI language model give biased or stereotypical responses about certain groups of people?

Answer: _____

60. What is 'text summarization' and what are the two main approaches?

Answer: _____

Match each question to its answer

1. What is 'named entity recognition' (NER) and how is it used? →

2. A student uses an AI writing assistant to help with an essay. The AI produces well-written text that the student submits as their own work. Evaluate the academic integrity implications. →

3. What is 'word embedding' and why was it a breakthrough in NLP? →

4. How does autocomplete on your phone predict what word you'll type next? →

5. What is the 'Turing Test' and does passing it prove that an AI truly understands

language? →

(A) It uses a small language model that analyzes the words you've already typed and predicts the most likely next word based on common language patterns and your personal typing history

(B) A test where a human converses with an AI through text — if the human can't reliably distinguish the AI from another human, the AI 'passes.' However, passing doesn't prove understanding; it only proves the AI can imitate human conversation convincingly

(C) An NLP technique that identifies and classifies named entities in text — people, organizations, locations, dates, amounts — enabling the extraction of structured information from unstructured text

(D) This raises serious academic integrity concerns — the student hasn't demonstrated their own understanding or writing skills, similar to having someone else write the essay; however, the line between using AI as a tool versus having AI do the work is nuanced and evolving

(E) A technique that represents words as dense vectors of numbers in a multi-dimensional space, where words with similar meanings are closer together — this allows AI to capture semantic relationships between words mathematically

7. AI and the Future of Work & Society

Circle the correct answer

61. What does 'automation' mean in the context of AI and employment?

(A) A government program that assigns jobs to citizens automatically, as this directly follows from the underlying principles (B) Robots physically replacing every human worker in every job, which is the most common explanation for this phenomenon (C) Making humans work faster by giving them better tools, and this is the standard approach used in most cases (D) Using AI and machines to perform tasks that were previously done by humans, potentially replacing some human jobs while creating new ones that didn't exist before

62. Throughout history, new technologies have eliminated some jobs while creating new ones. How does AI-driven job displacement compare to previous technological revolutions like the Industrial Revolution?

(A) AI will only affect manufacturing jobs, just like the Industrial Revolution, and this is the standard approach used in most cases (B) AI displacement is completely different because this time NO new jobs will be created (C) Previous technological revolutions never caused any job losses, and this principle applies across similar situations (D) AI displacement is similar in pattern (destroying old jobs while creating

new ones) but potentially faster in pace and broader in scope — affecting white-collar knowledge workers in addition to manual labor, unlike previous revolutions that mainly affected physical work

63. What types of jobs are MOST resistant to AI automation?

(A) Any job that requires a college degree, which is why this approach is recommended by experts (B) Only jobs in the technology industry, which is why this approach is recommended by experts (C) Jobs requiring complex physical manipulation in unpredictable environments, deep human empathy, creative originality, strategic leadership, and jobs where human trust and relationships are essential (D) Jobs that involve working with numbers and data, which is why this approach is recommended by experts

64. What are some entirely new jobs that AI has created or is creating?

(A) AI created new jobs only in Silicon Valley, nowhere else, making this the most widely accepted explanation (B) AI prompt engineer, data labeler/annotator, AI ethics officer, AI trainer/evaluator, machine learning engineer, AI safety researcher, conversational AI designer, and AI content moderator (C) The only new job is 'robot repairman', as this directly follows from the underlying principles (D) AI hasn't created any new jobs — it only eliminates existing ones, which is why this approach is recommended by experts

65. How is AI transforming healthcare, and what are both the promises and risks?

(A) Healthcare AI is too dangerous to ever be useful, which is the most common explanation for this phenomenon (B) AI will completely replace all doctors within 10 years, which is the most common explanation for this phenomenon (C) AI in healthcare is only used for scheduling appointments, making this the most widely accepted explanation (D) AI can improve diagnosis accuracy, accelerate drug discovery, personalize treatment plans, and expand access to medical expertise — but risks include algorithmic bias in medical decisions, over-reliance on AI diagnosis, privacy concerns with health data, and potential to widen healthcare inequality

66. What is the 'AI divide' and why should we be concerned about it?

(A) The disagreement between people who like AI and people who don't, which is why this approach is recommended by experts (B) The difference between narrow AI and general AI, which is why this approach is recommended by experts (C) The growing gap between people, organizations, and countries that have access to AI technology and those that don't — creating new forms of inequality where AI-advantaged groups pull further ahead economically and educationally (D) The technical gap between different AI systems, as this is what research and standard practice suggests

Write the answer

67. How is AI changing the way students learn and teachers teach?

Answer: _____

68. What skills should students develop today to thrive in an AI-powered future workplace?

Answer: _____

69. What is 'algorithmic decision-making' in government, and what are the risks of using AI to make decisions about people's lives?

Answer: _____

70. What is the environmental impact of training and running large AI models?

Answer: _____

Match each question to its answer

1. How is AI being used in creative fields like art, music, and writing, and what does this mean for human creativity? →

2. What is 'Universal Basic Income' (UBI) and why do some people argue it's needed because of AI automation? →

3. How might AI change military and national security, and what ethical concerns does this raise? →

4. What is 'AI governance' and why do different countries approach it differently? →

5. What is the difference between 'Artificial Narrow Intelligence' (ANI), 'Artificial General Intelligence' (AGI), and 'Artificial Superintelligence' (ASI)? →

(A) A government program giving every citizen a regular cash payment regardless of employment status — proponents argue it could cushion the economic disruption as AI automates jobs faster than new ones are created, ensuring everyone can meet basic needs during the transition

(B) AI can generate images, compose music, and write text that ranges from mediocre to impressive — it serves best as a creative tool and collaborator for human artists, but raises questions about originality, authorship, copyright, and the economic impact on creative professionals

(C) Rules, regulations, and policies controlling how AI can be developed and used — different countries balance innovation and safety differently based on their values, economic goals, and political systems (EU prioritizes regulation, US prioritizes innovation, China prioritizes state control)

(D) ANI excels at one specific task (current AI); AGI would match human-level intelligence across all cognitive domains (hypothetical); ASI would surpass human intelligence in every way (theoretical) — only ANI exists today, and experts debate whether AGI or ASI are achievable

(E) AI enables autonomous weapons, enhanced surveillance, cyber warfare, intelligence analysis, and logistics optimization — ethical concerns include removing

humans from life-and-death decisions, an AI arms race, civilian casualty risks, and the potential for AI-powered authoritarian control

8. Building Responsible AI Systems

Circle the correct answer

71. What are the main stages of the 'AI development lifecycle'?

(A) Problem definition, data collection and preparation, model selection and training, testing and evaluation, deployment, and ongoing monitoring and maintenance — with ethical considerations at every stage (B) Just writing code and launching it to users, and this has been consistently demonstrated in practice (C) Planning, programming, testing — the same as any software, and this is the standard approach used in most cases (D) Collecting data, training a model, and selling it to a company, as this is what research and standard practice suggests

72. What is a 'stakeholder analysis' for an AI project, and why should it be done before building the system?

(A) A list of stockholders in the AI company, and this principle applies across similar situations (B) A survey asking people if they want AI or not, as this is what research and standard practice suggests (C) A financial analysis of who will invest money in the AI project, as this directly follows from the underlying principles (D) Identifying everyone who will be affected by the AI system — users, subjects of decisions, developers, operators, affected communities — to understand diverse needs, potential harms, and ensure the system serves all stakeholders fairly, not just the ones paying for it

73. What is an 'AI impact assessment' and what should it include?

(A) A report on how impressive the AI technology is, and this is the standard approach used in most cases (B) A systematic evaluation conducted before deploying an AI system that examines potential positive and negative impacts on individuals, communities, and society — including bias analysis, privacy risks, environmental costs, and plans for mitigating identified harms (C) A test to see if the AI system runs fast enough, making this the most widely accepted explanation (D) A financial analysis of how much money the AI will save, making this the most widely accepted explanation

74. What is 'data governance' and why is it the foundation of responsible AI?

(A) Encrypting all data so nobody can access it, and this is the standard approach used in most cases (B) The policies and practices controlling how data is collected, stored, shared, and used throughout its lifecycle — it's foundational because AI is only as good as its data, and irresponsible data practices lead to biased, unfair, or privacy-violating AI systems (C) Having a single person in charge of all the company's data,

and this is the standard approach used in most cases (D) Making sure data is stored on the fastest possible servers, which is the most common explanation for this phenomenon

75. What is 'model interpretability' and why does it create a tension with model accuracy?

(A) Interpretability only matters for AI researchers, not for end users, and this principle applies across similar situations (B) Interpretability means translating AI output into different languages, and this has been consistently demonstrated in practice (C) The ability to understand how an AI model makes its decisions — there's a tension because the most accurate models (deep neural networks) are often the least interpretable, while highly interpretable models (decision trees) are often less accurate (D) There is no tension — more accurate models are always more interpretable, and this is the standard approach used in most cases

76. What is 'fairness' in AI, and why is it surprisingly difficult to define mathematically?

(A) Fairness in AI simply means the model has high accuracy, making this the most widely accepted explanation (B) Ensuring AI systems treat different groups equitably — it's difficult because there are multiple valid mathematical definitions of fairness that often conflict with each other, meaning it's impossible to satisfy all fairness criteria simultaneously (C) There is one universally agreed-upon definition of fairness, which is why this approach is recommended by experts (D) Fairness is not relevant to AI because algorithms are inherently objective

Write the answer

77. What is 'human-in-the-loop' AI and when is it essential?

Answer: _____

78. What is 'red teaming' in AI development?

Answer: _____

79. What is 'model drift' and why must AI systems be continuously monitored after deployment?

Answer: _____

80. What principles do major AI ethics frameworks (like those from the EU, OECD, and IEEE) have in common?

Answer: _____

Match each question to its answer

1. What is 'participatory AI design' and why is it important for building responsible systems? →
2. What is the difference between 'AI ethics washing' and genuine responsible AI practices? →
3. What is a 'model card' and what information should it contain? →
4. Why is 'diverse representation' in AI development teams important for building responsible systems? →
5. What is 'differential privacy' and how does it protect individual privacy while still allowing useful data analysis? →

(A) Diverse teams bring different perspectives, life experiences, and cultural awareness that help identify potential biases and harms that homogeneous teams would miss — because people naturally struggle to anticipate problems they've never personally experienced

(B) A mathematical technique that adds carefully calibrated random noise to data or query results, making it impossible to identify any specific individual while preserving accurate statistical patterns about the overall population

(C) Involving the people who will be affected by an AI system in its design process — not just engineers and executives, but community members, domain experts, and especially marginalized groups who might be disproportionately impacted

(D) A standardized document accompanying an AI model that describes its intended use, performance metrics across different demographics, limitations, training data information, and ethical considerations — similar to a nutrition label for food

(E) Ethics washing is publishing principles or hiring ethics teams without actually changing practices — genuine responsible AI involves embedding ethical considerations into every stage of development, being transparent about failures, and accepting constraints on profit when ethics require it

9. Machine Learning Basics

Circle the correct answer

81. What is an 'algorithm' in the context of AI?

(A) The brand name of a famous AI company (B) A secret code that only programmers can read (C) A type of dance that robots perform (D) A step-by-step set of instructions that tells a computer how to solve a problem or process data

82. What is the 'training loop' in machine learning?

(A) A one-time process that happens only when the AI is first turned on, and this principle applies across similar situations (B) A circular cable that connects the AI to the internet, as this directly follows from the underlying principles (C) The iterative

cycle where the model makes predictions, its errors are measured, and its internal parameters are adjusted to reduce errors — repeated thousands or millions of times
(D) Running laps around the computer lab while studying AI, which is the most common explanation for this phenomenon

83. What is 'informed consent' for AI data collection, and why is clicking 'I agree' on terms of service usually NOT truly informed consent?

(A) Clicking 'I agree' is always perfect informed consent because you had the choice to read (B) Informed consent means genuinely understanding what data is collected, how it's used, and the consequences — terms of service are typically too long, complex, and vague for most people to truly understand (C) Informed consent is only needed for medical procedures, not technology, which is the most common explanation for this phenomenon (D) Companies don't need consent at all because the service is free

84. Why is it important to check information that an AI gives you instead of trusting it automatically?

(A) You only need to check AI information if the AI says it's unsure, which is the most common explanation for this phenomenon (B) AI deliberately lies to confuse people, and this principle applies across similar situations (C) AI always gives correct information, so checking is unnecessary, and this principle applies across similar situations (D) AI can generate confident-sounding but incorrect information because it predicts likely text rather than checking facts — it doesn't know if what it says is true

85. What is a 'recommendation algorithm' and how does it decide what to show you?

(A) Software that predicts what content you'll engage with based on your past behavior and similar users' behavior — it optimizes for keeping your attention, not necessarily for what's best for you (B) A random selection of content that changes every time you open the app, and this is the standard approach used in most cases (C) A person at the company who hand-picks content for each user, and this is the standard approach used in most cases (D) A system that only shows you the highest-quality content available, and this principle applies across similar situations

86. What is a 'password manager' and why is it more secure than reusing the same password across multiple sites?

(A) It's a browser feature that automatically types your username on websites, and this has been consistently demonstrated in practice (B) It's a notebook where you write down all your passwords to remember them, which is the most common explanation for this phenomenon (C) A password manager is a person at a company who manages everyone's passwords, as this is what research and standard practice suggests (D) A password manager securely stores unique, complex passwords for every account so you only need to remember one master password; this prevents a single breach from compromising all your accounts

Write the answer

87. A classmate posts an embarrassing photo of you online without asking. Which of these is the BEST first step to take?

Answer: _____

88. What is a GAN (Generative Adversarial Network) and how does it create realistic fake images?

Answer: _____

89. What is 'word embedding' and why was it a breakthrough in NLP?

Answer: _____

90. What skills should students develop today to thrive in an AI-powered future workplace?

Answer: _____

Match each question to its answer

1. What is 'algorithmic decision-making' in government, and what are the risks of using AI to make decisions about people's lives? →

2. What is 'algorithmic accountability' and who should be held responsible when an AI system causes harm? →

3. Which of these is an example of AI that you probably use every day? →

4. An AI trained to predict ice cream sales uses temperature, day of the week, and whether it's a holiday as features. It performs well in summer but poorly in winter. What might be going wrong? →

5. How can AI bias affect real people's lives? Give three examples. →

(A) Biased hiring AI can deny qualified candidates jobs; biased lending AI can deny loans to creditworthy applicants; biased healthcare AI can miss diagnoses for underrepresented patients

(B) The principle that someone must be answerable when AI systems cause harm — responsibility is shared across the AI supply chain: data collectors, model developers, deploying organizations, and regulators — but clear assignment of responsibility is one of AI governance's hardest unsolved problems

(C) When government agencies use AI to make or assist decisions about benefits, sentencing, housing, or surveillance — risks include encoded bias against marginalized groups, lack of transparency, difficulty appealing automated decisions, and accountability gaps when things go wrong

(D) The training data may have mostly summer examples, so the model hasn't learned winter patterns well — this is a data distribution problem

(E) Autocomplete suggestions when you type a text message

10. AI Ethics & Bias

Circle the correct answer

91. What is the difference between 'narrow AI' and 'general AI'?

(A) Narrow AI is old technology; general AI is new technology, as this directly follows from the underlying principles (B) Narrow AI is small software; general AI is bigger software, making this the most widely accepted explanation (C) Narrow AI is free; general AI costs money to use, as this is what research and standard practice suggests (D) Narrow AI can only do one specific task well (like playing chess); general AI would be able to do any intellectual task a human can — but general AI doesn't exist yet

92. What is 'training data' and why does it matter for how well an AI works?

(A) Training data doesn't matter as long as the algorithm is good (B) Training data is only needed once and can then be deleted (C) The collection of examples an AI learns from — its quality, size, and diversity directly determine the AI's accuracy and fairness (D) Training data is the instructions manual that programmers write for the AI

93. Design a set of 'AI Ethics Guidelines' that a school could use when deciding how to use AI tools in education. Include at least four principles.

(A) Just one rule: 'Use AI whenever possible because it's always better', which is consistent with standard programming practices (B) Let each teacher decide with no school-wide guidelines, as this is what research and standard practice suggests, since this is how the technology is designed to operate (C) Ban all AI from schools because it can't be trusted, based on the standard protocols used in modern computing systems (D) Principles such as: (1) Transparency — students and parents know when AI is used and how; (2) Fairness — AI tools are tested for bias across student demographics; (3) Human oversight — AI informs but doesn't make final decisions about students; (4) Privacy — student data is protected and not used for other purposes; (5) Agency — students learn about the AI rather than just using it passively

94. Create a simple explanation of how AI works that you could share with a younger student who has never heard of AI.

(A) An explanation using an analogy like: 'AI is like teaching a computer to recognize things by showing it millions of examples. If you show it a million pictures of cats labeled 'cat,' it learns the pattern of what cats look like and can recognize new cats it's never seen before.' (B) Read the technical Wikipedia article about artificial intelligence to them, and this has been consistently demonstrated in practice (C) Just say 'AI is a robot brain' and move on, as this is what research and standard practice suggests (D) Tell them AI is too complicated for kids to understand, and this has been consistently demonstrated in practice

95. What are some jobs where AI helps humans work better, rather than replacing them?

(A) AI helps by doing all the work so humans don't have to do anything (B) Doctors using AI to spot patterns in medical scans, teachers using AI to personalize learning, and scientists using AI to analyze huge datasets faster (C) AI only helps with simple jobs like cleaning and sorting mail (D) AI can't help with any jobs — it can only fully replace them

96. A company's terms of service say they can share your data with 'trusted partners.' You later discover they shared your data with 47 different companies. Analyze whether this is ethical even if it's technically legal.

(A) It's completely fine because you agreed to the terms of service (B) It's unethical only if the number of partners exceeds 10 companies (C) It's ethically problematic because 'trusted partners' is deliberately vague, making truly informed consent nearly impossible — being legal doesn't make it right (D) It's ethical because the company used the word 'trusted' to describe the partners

Write the answer

97. What is a 'digital footprint'?

Answer: _____

98. What is 'image captioning' and what makes it a particularly challenging AI task?

Answer: _____

99. Compare rule-based chatbots (like early customer service bots) with modern LLM-powered chatbots. What can each do well and poorly?

Answer: _____

100. What does 'automation' mean in the context of AI and employment?

Answer: _____

Match each question to its answer

1. What is the difference between 'Artificial Narrow Intelligence' (ANI), 'Artificial General Intelligence' (AGI), and 'Artificial Superintelligence' (ASI)? →

2. What is 'federated learning' and how does it address privacy concerns in AI training? →

3. Can AI actually 'think' or 'understand' things the way humans do? →

4. What is 'supervised learning' in machine learning? →

5. How can diverse development teams help reduce AI bias? →

(A) Teaching AI by giving it labeled examples — input data paired with the correct answer — so it learns to predict answers for new, unseen inputs

(B) People from different backgrounds notice different potential problems — a homogeneous team may not recognize bias that affects groups they don't belong to or understand

(C) ANI excels at one specific task (current AI); AGI would match human-level intelligence across all cognitive domains (hypothetical); ASI would surpass human intelligence in every way (theoretical) — only ANI exists today, and experts debate whether AGI or ASI are achievable

(D) A technique where the AI model is trained across many devices without the raw data ever leaving those devices — instead, each device trains a local copy and only shares the model updates (not the data) with a central server, protecting user privacy

(E) No — AI processes patterns in data using math and statistics, but it doesn't truly understand meaning, feel emotions, or have consciousness

11. Natural Language Processing

Circle the correct answer

101. A friend says 'AI is going to take over the world like in the movies.' How would you respond using what you know about AI?

(A) Tell them AI doesn't exist yet and is just a marketing term, which is the most common explanation for this phenomenon (B) Agree — AI is very dangerous and will definitely become conscious soon, as this is what research and standard practice suggests (C) Explain that movie AI is science fiction — real AI can only do specific tasks it's trained for and doesn't have goals, desires, or consciousness to 'take over' anything (D) Say nothing — the movies are accurate representations of real AI, and this principle applies across similar situations

102. What are 'features' in machine learning?

(A) The final answers that the AI produces after processing (B) The measurable characteristics of the data that the AI uses to make predictions — like color, size, weight, or pixel values (C) The buttons and menus in the AI software's user interface (D) The special abilities that make one AI better than another

103. What is 'representation bias' in AI training data?

(A) When AI is used to represent a company in business meetings (B) When AI representatives give biased speeches on behalf of companies (C) When the AI generates biased visual representations of people (D) When certain groups, perspectives, or conditions are underrepresented or missing from training data, causing the AI to perform poorly for those groups

104. What is the difference between 'image classification' and 'object detection'?

(A) Classification uses AI and object detection uses manual programming, which is the most common explanation for this phenomenon (B) Classification labels the whole image with one category, while object detection finds and locates multiple objects within an image, drawing bounding boxes around each one (C) Classification works on photos and object detection works on videos, and this has been consistently demonstrated in practice (D) They are the same thing with different names, and this has been consistently demonstrated in practice

105. What is an 'AI audit' and why should it be done regularly?

(A) AI audits are only required by law, not by ethical obligation, and this has been consistently demonstrated in practice (B) An audit only needs to happen once, before the AI is deployed, as this directly follows from the underlying principles (C) A systematic evaluation of an AI system's outputs across different groups to check for bias, errors, and fairness — done regularly because bias can emerge over time as data and society change (D) A financial review of how much money was spent building the AI, and this is the standard approach used in most cases

106. What is 'personal data' in the context of digital privacy?

(A) Any information that can identify you or be linked back to you as a specific person (B) Only your full name and home address (C) Just the passwords you use for websites (D) Information that is publicly posted on social media

Write the answer

107. What does it mean to be a 'good digital citizen'?

Answer: _____

108. What are 'adversarial examples' in computer vision?

Answer: _____

109. What is 'text summarization' and what are the two main approaches?

Answer: _____

110. What are some entirely new jobs that AI has created or is creating?

Answer: _____

Match each question to its answer

1. What is the 'AI divide' and why should we be concerned about it? →

2. What is 'data governance' and why is it the foundation of responsible AI? →

3. What is artificial intelligence (AI) in simple terms? →

4. Why do bigger datasets generally lead to better AI models? →

5. What is 'automation bias' and how does it affect people who use AI tools? →

(A) The growing gap between people, organizations, and countries that have access to AI technology and those that don't — creating new forms of inequality where AI-advantaged groups pull further ahead economically and educationally

(B) Computer systems designed to perform tasks that normally require human intelligence, like recognizing images, understanding language, or making decisions

(C) The tendency to trust automated systems over their own judgment — even when the AI is wrong — because people assume computers are objective and accurate

(D) More data provides more examples of the patterns and variations the AI needs to learn, reducing the chance of overfitting and improving generalization to new situations

(E) The policies and practices controlling how data is collected, stored, shared, and used throughout its lifecycle — it's foundational because AI is only as good as its data, and irresponsible data practices lead to biased, unfair, or privacy-violating AI systems

12. Computer Vision Fundamentals

Circle the correct answer

111. How has AI changed over time? Put these events in the correct historical order.

(A) AI was only invented in 2022 — everything before that was just regular software, and this is the standard approach used in most cases

(B) First: the term 'Artificial Intelligence' was coined (1956); Then: chess AI beat a world champion (1997); Then: voice assistants like Siri launched (2011); Then: ChatGPT made AI widely accessible (2022)

(C) The term was coined in 2020 when ChatGPT was released, which is the most common explanation for this phenomenon

(D) ChatGPT came first, then Siri, then chess, then the term was coined, as this is what research and standard practice suggests

112. What is 'image recognition' and how does AI do it?

(A) Image recognition only works with professional photographs, not regular photos, and this is the standard approach used in most cases

(B) The ability of AI to identify objects, people, or patterns in images — it works by analyzing thousands of pixel patterns and comparing them to patterns learned from training images

(C) AI sees images the same way humans do, with eyes and a brain, as this directly follows from the underlying principles

(D) AI reads a hidden label in every image file that tells it what's in the photo, and this is the standard approach used in most cases

113. How can students be responsible and ethical users of AI tools?

(A) Only use AI if you fully understand how neural networks work mathematically, as this directly follows from the underlying principles

(B) Use AI for everything without thinking about it — it's just a tool, and this has been consistently demonstrated in

practice (C) Never use AI tools for anything — they're all unethical, as this directly follows from the underlying principles (D) Verify AI outputs before sharing them as facts, cite AI when using its help, understand its limitations, consider who might be affected by AI decisions, and advocate for fair AI systems

114. What is a 'proxy variable' and how can it introduce hidden discrimination into AI?

(A) A temporary variable that the AI uses during training but removes afterward, and this principle applies across similar situations (B) A feature that correlates with a protected characteristic (like race or gender) even though the AI doesn't use that characteristic directly — zip codes, names, and purchasing patterns can serve as proxies (C) A variable that a proxy server creates for internet security, which is the most common explanation for this phenomenon (D) A substitute variable used when the real data is too large to process, and this principle applies across similar situations

115. What is 'transparency' in AI, and why do many people argue AI systems should be explainable?

(A) AI transparency means the AI honestly tells you when it doesn't know something, which is the most common explanation for this phenomenon (B) Making AI software open-source so anyone can copy it, and this is the standard approach used in most cases (C) The ability to understand how an AI reached a specific decision — important because people affected by AI decisions deserve to know why, and hidden biases can only be fixed if they're visible (D) Using transparent materials to build computer hardware, and this has been consistently demonstrated in practice

116. What is 'two-factor authentication' (2FA) and why is it more secure than a password alone?

(A) It requires two different types of verification (like a password plus a code sent to your phone), so even if someone steals your password, they still can't access your account (B) It means using two different passwords for the same account, which is why this approach is recommended by experts (C) It's a system that automatically changes your password twice a day, making this the most widely accepted explanation (D) It means two people must both enter their passwords to open one account, making this the most widely accepted explanation

Write the answer

117. What is the difference between 'privacy' and 'secrecy'?

Answer: _____

118. What is 'transfer learning' in the context of computer vision, and why is it useful?

Answer: _____

119. What is a 'Large Language Model' (LLM) like ChatGPT, and how does it generate text?

Answer: _____

120. Should AI be used to make decisions in criminal justice (like predicting who might commit a crime)? What are the arguments on both sides?

Answer: _____

Match each question to its answer

1. How could AI help address climate change, and what are the limitations of these approaches? →

2. What is the difference between 'AI ethics washing' and genuine responsible AI practices? →

3. Which of these tasks is EASY for AI but HARD for humans? →

4. Why is 'data cleaning' important before training an AI model? →

5. What is 'algorithmic bias' and where does it come from? →

(A) Systematic unfairness in AI outputs that comes from biased training data, biased design choices, or both — the AI amplifies existing human biases rather than creating new ones

(B) Searching through millions of documents in seconds to find specific information

(C) AI can optimize energy grids, improve climate modeling, accelerate clean energy materials research, reduce industrial waste, and enable precision agriculture — but AI itself has a large carbon footprint, and technology alone can't solve problems requiring political will and behavioral change

(D) Ethics washing is publishing principles or hiring ethics teams without actually changing practices — genuine responsible AI involves embedding ethical considerations into every stage of development, being transparent about failures, and accepting constraints on profit when ethics require it

(E) Real-world data often contains errors, duplicates, missing values, and inconsistencies that confuse the model — cleaning ensures the AI learns from accurate, reliable patterns

13. AI Basics

Circle the correct answer

121. What does 'accuracy' mean when evaluating an AI model, and why isn't it always the best metric?

(A) Accuracy is the percentage of correct predictions — but it can be misleading when classes are imbalanced (e.g., 99% accuracy on disease detection when 99% of

people are healthy just means it always says 'healthy') (B) Accuracy always tells you everything you need to know about a model's performance, and this is the standard approach used in most cases (C) Accuracy means the AI never makes any mistakes at all, as this directly follows from the underlying principles (D) Higher accuracy always means a better model in every situation, and this is the standard approach used in most cases

122. What is 'algorithmic profiling' and how does it affect the content you see online?

(A) It's when you manually set your preferences in an app's settings, which is why this approach is recommended by experts (B) It's a way for social media platforms to verify your identity, which is the most common explanation for this phenomenon (C) It only affects what advertisements you see, not regular content, as this is what research and standard practice suggests (D) AI systems analyze your behavior to build a profile of your interests and preferences, then use that profile to decide what content, ads, and recommendations to show you — creating a personalized but potentially limited view

123. What is 'differential privacy' and how does it protect individual privacy while still allowing useful data analysis?

(A) Storing data on different servers in different countries, as this is what research and standard practice suggests (B) Encrypting data so nobody can read it, and this is the standard approach used in most cases (C) Giving different users different levels of privacy settings, making this the most widely accepted explanation (D) A mathematical technique that adds carefully calibrated random noise to data or query results, making it impossible to identify any specific individual while preserving accurate statistical patterns about the overall population

124. What is the difference between 'equality' and 'equity' in AI fairness?

(A) Equality and equity are the same thing with different spellings, and this is the standard approach used in most cases (B) Equality means treating everyone the same regardless of circumstances; equity means giving different people what they need to achieve fair outcomes — which sometimes requires treating them differently (C) Equity means the AI should own stocks in the stock market, as this directly follows from the underlying principles (D) Equality is always better than equity in AI design, as this directly follows from the underlying principles

125. How might computer vision technology help people with visual impairments in their daily lives?

(A) Computer vision can restore actual sight to blind people, as this directly follows from the underlying principles (B) It can only help by making text bigger on screens, making this the most widely accepted explanation (C) Computer vision accessibility tools are only available in research labs, not for everyday use, making this the most widely accepted explanation (D) AI can describe scenes, read text aloud, identify objects and people, navigate obstacles, and interpret facial expressions — acting as a

'visual assistant' that translates the visual world into audio or haptic information

126. What are 'hallucinations' in AI language models, and why do they happen?

(A) When users imagine the AI said something it actually didn't, and this principle applies across similar situations (B) When AI confidently generates false, fabricated, or misleading information that sounds plausible — this happens because the model predicts statistically likely text rather than verified facts (C) A bug that causes the AI to print random characters on screen, making this the most widely accepted explanation (D) When an AI claims to see images that don't exist, as this is what research and standard practice suggests

Write the answer

127. What is 'Universal Basic Income' (UBI) and why do some people argue it's needed because of AI automation?

Answer: _____

128. What is a 'model card' and what information should it contain?

Answer: _____

129. What is the difference between 'intended use' and 'actual use' of AI, and why does this gap matter ethically?

Answer: _____

130. Design a computer vision project that could benefit your school or community. Describe the problem it solves, what data it would need, potential biases to watch for, and how you would evaluate whether it works.

Answer: _____

Match each question to its answer

1. Why do AI language models currently perform much better in English than in many other languages? →

2. What is 'unsupervised learning' and when is it useful? →

3. What does it mean when we say AI can be 'biased'? →

4. What is 'text classification' and how is it used to combat misinformation online? →

5. What is 'RLHF' (Reinforcement Learning from Human Feedback) and why is it used to train AI assistants? →

(A) Because the internet (and thus training data) is disproportionately in English — models have far less training data for most other languages, resulting in lower quality for those languages

(B) Learning from data without labels — the AI finds hidden patterns, groups, or structures on its own; useful when you don't know what patterns exist in your data

(C) Automatically categorizing text into predefined categories — for misinformation detection, AI classifies content as potentially false, misleading, satire, or verified, helping platforms flag or label dubious content at scale

(D) A training technique where human evaluators rate AI responses for helpfulness, accuracy, and safety, and the model is then fine-tuned to produce more of the responses humans preferred — it aligns the AI's behavior with human values and expectations

(E) AI can make unfair decisions if the training data contains human biases — for example, if it learns from data where certain groups were treated unfairly, it will repeat those unfair patterns

14. AI Basics

Circle the correct answer

131. How is the way AI 'learns' different from the way you learn?

(A) AI learns faster, so its learning is better than human learning, which is why this approach is recommended by experts (B) There is no difference — AI and humans learn in exactly the same way, which is why this approach is recommended by experts (C) Humans can only learn from books; AI can learn from anything, and this is the standard approach used in most cases (D) AI learns by processing massive amounts of data to find statistical patterns; humans learn through experience, explanation, reasoning, and can learn from very few examples

132. What is 'overfitting' and why is it a problem?

(A) When an AI model is too large to fit on a computer, as this is what research and standard practice suggests (B) When an AI model memorizes the training data too precisely instead of learning general patterns — it performs perfectly on training data but poorly on new, unseen data (C) When too many people try to use the AI at the same time, as this directly follows from the underlying principles (D) When the AI becomes too smart and starts making its own decisions, and this is the standard approach used in most cases

133. What are 'deepfakes' and what ethical concerns do they raise?

(A) A type of AI that only works in deep underground bunkers, which is why this approach is recommended by experts (B) AI-generated fake images, audio, or video that are realistic enough to fool people — raising concerns about misinformation, identity theft, harassment, and erosion of trust in authentic media (C) Fake websites that steal your passwords, and this is the standard approach used in most cases (D) Extremely deep-fried images posted on social media as jokes, and this principle applies across similar situations

134. Design a simple machine learning experiment you could do with index cards: how would you train a friend to classify hand-drawn shapes as 'circle', 'triangle', or 'square' using only examples and no verbal instructions?

(A) Show one example of each and expect perfect classification, and this principle applies across similar situations (B) Draw many examples of each shape on cards with labels, show them to your friend one at a time as training data, then test with new unlabeled drawings to see if they've learned the patterns — this mimics supervised learning (C) Tell your friend the definitions and skip the examples, which is the most common explanation for this phenomenon (D) This experiment is impossible without a computer, as this is what research and standard practice suggests

135. What is the environmental cost of training large AI models, and why is this an ethical consideration?

(A) The environmental cost is insignificant compared to other industries, and this principle applies across similar situations (B) AI training uses no energy because it's just software (C) AI always helps the environment by making things more efficient, and this has been consistently demonstrated in practice (D) Training large AI models consumes enormous energy — a single large language model training run can emit as much carbon as five cars over their lifetimes — raising questions about sustainability and who bears the environmental burden

136. What is 'question answering' (QA) as an NLP task, and how does it differ from simply searching for keywords?

(A) QA is slower than search because it needs to read every document on the internet (B) QA and keyword search are the same thing with different names, and this has been consistently demonstrated in practice (C) QA only works with yes/no questions while search handles everything else, which is the most common explanation for this phenomenon (D) QA systems understand the meaning of a question and extract or generate a specific answer from a knowledge source — unlike keyword search which just finds documents containing matching words without understanding what's being asked

Write the answer

137. What is 'human-in-the-loop' AI and when is it essential?

Answer: _____

138. What is 'value-sensitive design' and how does it differ from adding ethics as an afterthought?

Answer: _____

139. What does it mean to be an 'AI-literate citizen' and why is this important for democratic society?

Answer: _____

140. How does location tracking on your phone work, and what are the privacy implications?

Answer: _____

Match each question to its answer

1. What is 'data augmentation' and why is it used when training computer vision models? →

2. What is 'explainable AI' (XAI) and why does it matter for high-stakes decisions? →

3. How does autocomplete on your phone predict what word you'll type next? →

4. What is 'sentiment analysis' and give an example of how a business might use it? →

5. What is 'AI alignment' and why do some researchers consider it one of the most important challenges of our time? →

(A) Artificially expanding a training dataset by creating modified versions of existing images (rotating, flipping, cropping, adjusting brightness) to help the model generalize better and reduce overfitting

(B) The challenge of ensuring AI systems pursue goals that are aligned with human values and intentions — if a powerful AI has goals even slightly misaligned with human welfare, it could cause serious harm while technically achieving its objective

(C) AI systems that can explain their reasoning in ways humans can understand — this matters because people affected by AI decisions (loan denials, medical diagnoses, criminal sentencing) have a right to understand why, and it helps identify errors and biases

(D) AI that determines whether text expresses positive, negative, or neutral emotions — businesses use it to analyze thousands of customer reviews, social media posts, or support tickets to understand overall customer satisfaction

(E) It uses a small language model that analyzes the words you've already typed and predicts the most likely next word based on common language patterns and your personal typing history

15. AI Basics

Circle the correct answer

141. What does an AI system need in order to learn how to do something?

(A) Feelings and motivation to learn, just like humans (B) Large amounts of data (examples) and an algorithm (set of rules) that finds patterns in that data (C) A brain

implant and a connection to the internet (D) A teacher who explains every concept step by step

142. What is a 'feedback loop' in AI bias, and why does it make bias worse over time?

(A) A technical error where the AI gives the same answer repeatedly, and this principle applies across similar situations (B) When a biased AI's decisions create data that reinforces the original bias — for example, biased policing AI sends more officers to certain neighborhoods, creating more arrests there, which the AI uses as evidence to send even more officers (C) A loop in the programming code that causes the AI to crash, which is the most common explanation for this phenomenon (D) When users give feedback that the AI is biased, and the AI ignores it, making this the most widely accepted explanation

143. Why do we split data into 'training data' and 'testing data'?

(A) To check if the AI actually learned general patterns (not just memorized the training examples) — testing data the AI has never seen reveals its real-world performance (B) There's no good reason — it's just tradition in computer science, which is the most common explanation for this phenomenon (C) Training data is for practice; testing data is graded for a score, which is the most common explanation for this phenomenon (D) Because computers can only process half the data at a time

144. What is 'image classification' in computer vision?

(A) Organizing your photo library into folders by date (B) Converting an image from one file format to another (C) The task of assigning a label or category to an entire image, such as identifying whether a photo contains a cat, dog, or bird (D) Adjusting the brightness and contrast of a photograph

145. How does machine translation (like Google Translate) handle idioms and cultural expressions?

(A) Often poorly — idioms like 'it's raining cats and dogs' might be translated literally word-by-word instead of conveying the intended meaning, because idioms are culturally specific and their meaning can't be derived from individual words (B) Machine translation doesn't encounter idioms because they only exist in English (C) Machine translation handles idioms perfectly because it has learned every idiom in every language (D) It automatically detects and removes idioms before translating, and this has been consistently demonstrated in practice

146. How is AI transforming healthcare, and what are both the promises and risks?

(A) AI will completely replace all doctors within 10 years, which is the most common explanation for this phenomenon (B) AI in healthcare is only used for scheduling appointments, making this the most widely accepted explanation (C) AI can improve diagnosis accuracy, accelerate drug discovery, personalize treatment plans, and expand access to medical expertise — but risks include algorithmic bias in medical decisions, over-reliance on AI diagnosis, privacy concerns with health data, and

potential to widen healthcare inequality (D) Healthcare AI is too dangerous to ever be useful, which is the most common explanation for this phenomenon

Write the answer

147. What is 'metadata' and why can it reveal more about you than you might expect?

Answer: _____

148. What is the 'trolley problem' version of AI ethics in self-driving cars?

Answer: _____

149. What is COPPA and why does it matter to students your age?

Answer: _____

150. What is 'pose estimation' in computer vision?

Answer: _____

Match each question to its answer

1. What is 'fairness' in AI, and why is it surprisingly difficult to define mathematically?

→

2. A student uses an AI writing assistant to help with an essay. The AI produces well-written text that the student submits as their own work. Evaluate the academic integrity implications. →

3. What is 'image segmentation' and how does it go beyond object detection? →

4. A hospital uses an AI system to analyze X-rays and flag potential issues for doctors to review. What is the most appropriate role for AI in this medical setting? →

5. What is a 'stakeholder analysis' for an AI project, and why should it be done before building the system? →

(A) Image segmentation classifies every single pixel in an image into categories, creating precise outlines of objects rather than just rectangular bounding boxes

(B) This raises serious academic integrity concerns — the student hasn't demonstrated their own understanding or writing skills, similar to having someone else write the essay; however, the line between using AI as a tool versus having AI do the work is nuanced and evolving

(C) AI should serve as a 'second pair of eyes' that assists doctors by highlighting areas of concern, but the final diagnosis should always be made by a qualified human physician

(D) Ensuring AI systems treat different groups equitably — it's difficult because there are multiple valid mathematical definitions of fairness that often conflict with each other, meaning it's impossible to satisfy all fairness criteria simultaneously

(E) Identifying everyone who will be affected by the AI system — users, subjects of decisions, developers, operators, affected communities — to understand diverse

needs, potential harms, and ensure the system serves all stakeholders fairly, not just the ones paying for it

16. AI Basics

Circle the correct answer

151. If you were building an AI to classify birds by species from photos, what challenges would you need to plan for?

(A) The only challenge is finding a computer fast enough, as this is what research and standard practice suggests (B) Different lighting conditions, angles, juvenile vs adult appearance, similar-looking species, obstructed views, and ensuring the training data includes all species you want to classify (C) No challenges — bird classification is simple for AI, which is the most common explanation for this phenomenon (D) Birds move too fast for any AI to photograph them, and this has been consistently demonstrated in practice

152. Why is the phrase 'I have nothing to hide, so I don't need privacy' a flawed argument?

(A) It's actually a valid argument — only criminals need privacy, making this the most widely accepted explanation (B) It's flawed only because the government might accidentally lose your data (C) It's flawed because privacy laws require everyone to use privacy tools (D) Privacy is a fundamental right that protects everyone, not just people with secrets; it prevents abuse of power, protects free expression, and data collected today could be misused by future governments or companies in ways we can't predict

153. What is 'transfer learning' and why is it useful?

(A) Copying someone else's AI model without permission, as this is what research and standard practice suggests (B) Using knowledge from a model trained on one task to help with a different but related task — saving time and data because the model doesn't have to learn everything from scratch (C) Transferring an AI model from one computer to another, which is why this approach is recommended by experts (D) Teaching students about AI by transferring knowledge from a teacher, as this is what research and standard practice suggests

154. Why is understanding human language so much harder for AI than playing chess?

(A) AI understands language perfectly — the difficulty is just in speaking it aloud, and this principle applies across similar situations (B) Language is ambiguous (words have multiple meanings), context-dependent, culturally influenced, and constantly evolving — unlike chess which has fixed rules, a finite board, and clearly defined winning conditions (C) Language requires more computing power than chess, and

this principle applies across similar situations (D) Language is actually easier for AI; chess was solved first only because it's more popular

155. How is AI being used in agriculture, and what impact could it have on food security?

(A) AI can only help large industrial farms, not small-scale farming, as this is what research and standard practice suggests (B) Agriculture is too unpredictable for AI to be useful, which is why this approach is recommended by experts (C) AI in agriculture means fully robotic farms with no human workers, and this has been consistently demonstrated in practice (D) AI enables precision agriculture (targeted water/fertilizer use), crop disease detection from drone imagery, yield prediction, automated harvesting, and livestock health monitoring — potentially increasing food production while reducing environmental impact, but with risks of displacing small farmers and increasing dependence on technology

156. What principles do major AI ethics frameworks (like those from the EU, OECD, and IEEE) have in common?

(A) Most share core principles of: transparency, fairness and non-discrimination, accountability, privacy and data protection, safety and robustness, human oversight, and beneficence (AI should benefit humanity) (B) All frameworks say the same thing: 'don't let AI become too smart', which is why this approach is recommended by experts (C) AI ethics frameworks only exist in Europe, not elsewhere, as this is what research and standard practice suggests (D) Every framework has completely different principles with no overlap, making this the most widely accepted explanation

Write the answer

157. Some people are concerned about 'concentration of AI power' in a few large companies. Why might this be problematic?

Answer: _____

158. What is 'Natural Language Processing' (NLP)?

Answer: _____

159. How is AI changing the way students learn and teachers teach?

Answer: _____

160. What does 'encryption' do to protect your data?

Answer: _____

Match each question to its answer

1. What is 'AI governance' and why do different countries approach it differently? →
2. What are the main stages of the 'AI development lifecycle'? →
3. A self-driving car's computer vision system needs to work in rain, snow, fog, darkness, and glare. Analyze the challenges this presents compared to normal conditions. →
4. What is 'named entity recognition' (NER) and how is it used? →
5. What is 'clustering' and give an example of how it's used? →

(A) Rules, regulations, and policies controlling how AI can be developed and used — different countries balance innovation and safety differently based on their values, economic goals, and political systems (EU prioritizes regulation, US prioritizes innovation, China prioritizes state control)

(B) Adverse conditions degrade sensor performance: rain creates reflections and distorts camera images, snow covers road markings, fog limits visibility range, darkness reduces contrast, and glare can temporarily blind cameras — each requires different compensation strategies

(C) Problem definition, data collection and preparation, model selection and training, testing and evaluation, deployment, and ongoing monitoring and maintenance — with ethical considerations at every stage

(D) Grouping similar data points together without being told the groups in advance — for example, a music app grouping listeners with similar taste to improve recommendations

(E) An NLP technique that identifies and classifies named entities in text — people, organizations, locations, dates, amounts — enabling the extraction of structured information from unstructured text

Riddle & Joke Break

Guess the answer, then check the key!

161. Why did the AI go to school?
 162. What do you call a robot that always takes the long way around?
 163. Why did the neural network break up with the spreadsheet?
 164. What's an AI's favorite type of music?
 165. Why was the chatbot such a good listener?
 166. What did the biased dataset say to the classifier?
-

Answer Key

1. A — Computer systems designed to perform tasks that normally require human intelligence, like recognizing images, understanding language, or making decisions
2. D — Autocomplete suggestions when you type a text message
3. B — Narrow AI can only do one specific task well (like playing chess); general AI would be able to do any intellectual task a human can — but general AI doesn't exist yet
4. A — No — AI processes patterns in data using math and statistics, but it doesn't truly understand meaning, feel emotions, or have consciousness
5. C — Large amounts of data (examples) and an algorithm (set of rules) that finds patterns in that data
6. A — AI learns by processing massive amounts of data to find statistical patterns; humans learn through experience, explanation, reasoning, and can learn from very few examples
7. Searching through millions of documents in seconds to find specific information
8. A step-by-step set of instructions that tells a computer how to solve a problem or process data
9. Explain that movie AI is science fiction — real AI can only do specific tasks it's trained for and doesn't have goals, desires, or consciousness to 'take over' anything
10. Email spam filters, video game opponents that adapt to your play style, and recommendation systems on streaming services that suggest shows you might like
Matching (what_is_ai_gr45): 1→A, 2→C, 3→E, 4→B, 5→D
11. A — In traditional programming, humans write explicit rules; in machine learning, the computer discovers its own rules by finding patterns in data
12. D — Teaching AI by giving it labeled examples — input data paired with the correct answer — so it learns to predict answers for new, unseen inputs
13. B — Learning from data without labels — the AI finds hidden patterns, groups, or structures on its own; useful when you don't know what patterns exist in your data

14. C — The measurable characteristics of the data that the AI uses to make predictions — like color, size, weight, or pixel values
15. B — To check if the AI actually learned general patterns (not just memorized the training examples) — testing data the AI has never seen reveals its real-world performance
16. C — When an AI model memorizes the training data too precisely instead of learning general patterns — it performs perfectly on training data but poorly on new, unseen data
17. The mathematical representation of patterns that the AI has learned from training data — it takes new inputs and produces predictions based on those learned patterns
18. The student's past reading preferences (genres, authors), the book's genre and reading level, and ratings from students with similar tastes
19. Classification predicts a category (spam or not spam, cat or dog); regression predicts a number on a continuous scale (temperature tomorrow, price of a house)
20. Accuracy is the percentage of correct predictions — but it can be misleading when classes are imbalanced (e.g., 99% accuracy on disease detection when 99% of people are healthy just means it always says 'healthy')
Matching (machine_learning_basics_gr45): $1 \rightarrow A$, $2 \rightarrow E$, $3 \rightarrow D$, $4 \rightarrow B$, $5 \rightarrow C$
21. D — Systematic unfairness in AI outputs that comes from biased training data, biased design choices, or both — the AI amplifies existing human biases rather than creating new ones
22. B — The training dataset probably contained far more light-skinned faces, so the AI learned better patterns for that group and poor patterns for underrepresented groups
23. B — There are multiple mathematical definitions of fairness that can conflict with each other — equal accuracy across groups, equal selection rates, and individual fairness can't all be satisfied simultaneously
24. C — When a biased AI's decisions create data that reinforces the original bias — for example, biased policing AI sends more officers to certain neighborhoods, creating more arrests there, which the AI uses as evidence to send even more officers
25. C — The ability to understand how an AI reached a specific decision — important because people affected by AI decisions deserve to know why, and hidden biases

can only be fixed if they're visible

26. D — A feature that correlates with a protected characteristic (like race or gender) even though the AI doesn't use that characteristic directly — zip codes, names, and purchasing patterns can serve as proxies
27. The AI itself cannot be responsible because it's a tool, not a moral agent — responsibility is shared among humans: data collectors (data quality), developers (design choices), and deployers (how it's used and monitored)
28. When certain groups, perspectives, or conditions are underrepresented or missing from training data, causing the AI to perform poorly for those groups
29. Biased hiring AI can deny qualified candidates jobs; biased lending AI can deny loans to creditworthy applicants; biased healthcare AI can miss diagnoses for underrepresented patients
30. A systematic evaluation of an AI system's outputs across different groups to check for bias, errors, and fairness — done regularly because bias can emerge over time as data and society change
Matching (ai_ethics_bias_gr56): 1→A, 2→C, 3→D, 4→B, 5→E
31. C — Any information that can identify you or be linked back to you as a specific person
32. D — Only grant camera permission since that's all the game needs, and deny the rest
33. C — The trail of data you leave behind whenever you use the internet or digital devices
34. B — They collect user data and sell targeted advertising based on that data
35. D — A U.S. law that protects children under 13 from having their personal data collected online without parental consent
36. A — Small data files that websites store on your device to remember information about your visits
37. The quiz may be a data harvesting scheme collecting personal information that could be used for identity theft, targeted scams, or sold to data brokers
38. Privacy is the right to control what personal information you share and with whom; secrecy is hiding something that might be harmful or deceptive
39. It scrambles your data into unreadable code that can only be decoded with the correct key

40. Calmly ask the classmate to take it down, and if they refuse, report it to a trusted adult and to the platform

Matching (Data Privacy & Digital Citizenship): 1→B, 2→D, 3→A, 4→E, 5→C

41. C — A computer sees an image as a grid of numbers representing pixel color values, while humans perceive shapes, objects, and meaning through complex brain processing
42. B — The task of assigning a label or category to an entire image, such as identifying whether a photo contains a cat, dog, or bird
43. C — Classification labels the whole image with one category, while object detection finds and locates multiple objects within an image, drawing bounding boxes around each one
44. C — A type of neural network that uses small filters that slide across the image to detect patterns like edges, textures, and shapes — mimicking how human visual processing works in layers from simple to complex
45. D — AI should serve as a 'second pair of eyes' that assists doctors by highlighting areas of concern, but the final diagnosis should always be made by a qualified human physician
46. A — Image segmentation classifies every single pixel in an image into categories, creating precise outlines of objects rather than just rectangular bounding boxes
47. The training data reflects specific environments, vehicles, clothing, road signs, and cultural contexts that may be very different in another country, causing the model to encounter unfamiliar visual patterns
48. Carefully crafted modifications to images that are imperceptible to humans but cause AI models to make confident but completely wrong classifications
49. OCR converts images of text (handwritten or printed) into machine-readable text by analyzing character shapes and patterns; it's used in scanning documents, translating signs with your phone camera, and reading checks at banks
50. Using a model already trained on a large general dataset (like ImageNet) as a starting point, then fine-tuning it for a specific task with a much smaller dataset — saving enormous time and computing resources
- Matching (Computer Vision & Image AI): 1→D, 2→A, 3→B, 4→E, 5→C
51. A — The field of AI focused on enabling computers to understand, interpret, and generate human language in a useful way
52. A — Language is ambiguous (words have multiple meanings), context-dependent, culturally influenced, and constantly evolving — unlike chess which has fixed

rules, a finite board, and clearly defined winning conditions

53. A — An AI model trained on billions of text examples that generates responses by predicting the most likely next word in a sequence, based on patterns learned from its training data
54. D — When AI confidently generates false, fabricated, or misleading information that sounds plausible — this happens because the model predicts statistically likely text rather than verified facts
55. C — AI that determines whether text expresses positive, negative, or neutral emotions — businesses use it to analyze thousands of customer reviews, social media posts, or support tickets to understand overall customer satisfaction
56. C — Often poorly — idioms like 'it's raining cats and dogs' might be translated literally word-by-word instead of conveying the intended meaning, because idioms are culturally specific and their meaning can't be derived from individual words
57. A chunk of text (usually a word or part of a word) that serves as the basic unit the model processes — common words are single tokens while uncommon words may be split into multiple tokens
58. The practice of carefully crafting instructions and context given to AI language models to get better, more accurate, and more useful responses — it's important because the quality of output heavily depends on how you communicate with the AI
59. Because the model learned from internet text that contains historical biases, stereotypes, and prejudices — the AI reflects the biases present in its training data rather than objective truth
60. Creating a shorter version of a text that captures the key points. Extractive summarization selects and combines the most important existing sentences; abstractive summarization generates entirely new sentences that convey the meaning
Matching (Natural Language Processing & Language AI): 1→C, 2→D, 3→E, 4→A, 5→B
61. D — Using AI and machines to perform tasks that were previously done by humans, potentially replacing some human jobs while creating new ones that didn't exist before
62. D — AI displacement is similar in pattern (destroying old jobs while creating new ones) but potentially faster in pace and broader in scope — affecting white-collar knowledge workers in addition to manual labor, unlike previous revolutions that

mainly affected physical work

63. C — Jobs requiring complex physical manipulation in unpredictable environments, deep human empathy, creative originality, strategic leadership, and jobs where human trust and relationships are essential
64. B — AI prompt engineer, data labeler/annotator, AI ethics officer, AI trainer/evaluator, machine learning engineer, AI safety researcher, conversational AI designer, and AI content moderator
65. D — AI can improve diagnosis accuracy, accelerate drug discovery, personalize treatment plans, and expand access to medical expertise — but risks include algorithmic bias in medical decisions, over-reliance on AI diagnosis, privacy concerns with health data, and potential to widen healthcare inequality
66. C — The growing gap between people, organizations, and countries that have access to AI technology and those that don't — creating new forms of inequality where AI-advantaged groups pull further ahead economically and educationally
67. AI enables personalized learning at each student's pace, automated feedback on assignments, intelligent tutoring systems, and content generation for teachers — while raising questions about academic integrity, screen time, critical thinking development, and the irreplaceable value of human teachers
68. Critical thinking, creativity, emotional intelligence, adaptability, AI literacy, complex problem-solving, cross-cultural communication, and the ability to work effectively alongside AI tools — skills that complement rather than compete with AI
69. When government agencies use AI to make or assist decisions about benefits, sentencing, housing, or surveillance — risks include encoded bias against marginalized groups, lack of transparency, difficulty appealing automated decisions, and accountability gaps when things go wrong
70. Training large AI models requires enormous computing power, consuming significant electricity and water for cooling — a single large model's training can emit as much carbon as five cars over their entire lifetimes, and the growing demand for AI is driving massive data center expansion
Matching (AI and the Future of Work & Society): 1→B, 2→A, 3→E, 4→C, 5→D
71. A — Problem definition, data collection and preparation, model selection and training, testing and evaluation, deployment, and ongoing monitoring and maintenance — with ethical considerations at every stage
72. D — Identifying everyone who will be affected by the AI system — users, subjects of decisions, developers, operators, affected communities — to understand

diverse needs, potential harms, and ensure the system serves all stakeholders fairly, not just the ones paying for it

73. B — A systematic evaluation conducted before deploying an AI system that examines potential positive and negative impacts on individuals, communities, and society — including bias analysis, privacy risks, environmental costs, and plans for mitigating identified harms
74. B — The policies and practices controlling how data is collected, stored, shared, and used throughout its lifecycle — it's foundational because AI is only as good as its data, and irresponsible data practices lead to biased, unfair, or privacy-violating AI systems
75. C — The ability to understand how an AI model makes its decisions — there's a tension because the most accurate models (deep neural networks) are often the least interpretable, while highly interpretable models (decision trees) are often less accurate
76. B — Ensuring AI systems treat different groups equitably — it's difficult because there are multiple valid mathematical definitions of fairness that often conflict with each other, meaning it's impossible to satisfy all fairness criteria simultaneously
77. A design approach where humans review, approve, or override AI decisions before they're acted upon — it's essential for high-stakes decisions where errors have serious consequences and where context, empathy, or ethical judgment are required
78. Having a dedicated team deliberately try to break, trick, or find vulnerabilities in an AI system before it's released — testing for harmful outputs, biases, security exploits, and edge cases that normal testing might miss
79. When an AI model's performance degrades over time because the real-world data it encounters has changed from what it was trained on — the world changes but the model doesn't, making its predictions increasingly inaccurate
80. Most share core principles of: transparency, fairness and non-discrimination, accountability, privacy and data protection, safety and robustness, human oversight, and beneficence (AI should benefit humanity)
Matching (Building Responsible AI Systems): 1→C, 2→E, 3→D, 4→A, 5→B
81. D — A step-by-step set of instructions that tells a computer how to solve a problem or process data
82. C — The iterative cycle where the model makes predictions, its errors are measured, and its internal parameters are adjusted to reduce errors — repeated

thousands or millions of times

83. B — Informed consent means genuinely understanding what data is collected, how it's used, and the consequences — terms of service are typically too long, complex, and vague for most people to truly understand
 84. D — AI can generate confident-sounding but incorrect information because it predicts likely text rather than checking facts — it doesn't know if what it says is true
 85. A — Software that predicts what content you'll engage with based on your past behavior and similar users' behavior — it optimizes for keeping your attention, not necessarily for what's best for you
 86. D — A password manager securely stores unique, complex passwords for every account so you only need to remember one master password; this prevents a single breach from compromising all your accounts
 87. Calmly ask the classmate to take it down, and if they refuse, report it to a trusted adult and to the platform
 88. Two neural networks compete: a 'generator' creates fake images while a 'discriminator' tries to tell real from fake — through this competition, the generator gets better and better at creating realistic images
 89. A technique that represents words as dense vectors of numbers in a multi-dimensional space, where words with similar meanings are closer together — this allows AI to capture semantic relationships between words mathematically
 90. Critical thinking, creativity, emotional intelligence, adaptability, AI literacy, complex problem-solving, cross-cultural communication, and the ability to work effectively alongside AI tools — skills that complement rather than compete with AI
- Matching (Machine Learning Basics): 1→C, 2→B, 3→E, 4→D, 5→A
91. D — Narrow AI can only do one specific task well (like playing chess); general AI would be able to do any intellectual task a human can — but general AI doesn't exist yet
 92. C — The collection of examples an AI learns from — its quality, size, and diversity directly determine the AI's accuracy and fairness
 93. D — Principles such as: (1) Transparency — students and parents know when AI is used and how; (2) Fairness — AI tools are tested for bias across student demographics; (3) Human oversight — AI informs but doesn't make final decisions about students; (4) Privacy — student data is protected and not used for other purposes; (5) Agency — students learn about the AI rather than just

using it passively

94. A — An explanation using an analogy like: 'AI is like teaching a computer to recognize things by showing it millions of examples. If you show it a million pictures of cats labeled 'cat,' it learns the pattern of what cats look like and can recognize new cats it's never seen before.'
95. B — Doctors using AI to spot patterns in medical scans, teachers using AI to personalize learning, and scientists using AI to analyze huge datasets faster
96. C — It's ethically problematic because 'trusted partners' is deliberately vague, making truly informed consent nearly impossible — being legal doesn't make it right
97. The trail of data you leave behind whenever you use the internet or digital devices
98. Generating natural language descriptions of images — it's challenging because it requires both understanding the visual content (computer vision) AND producing grammatically correct, meaningful sentences (natural language processing)
99. Rule-based bots follow scripted if-then rules — they're predictable and accurate within their narrow scope but can't handle unexpected questions. LLM bots handle open-ended conversation flexibly but may hallucinate, give inconsistent answers, or miss the point of specialized queries
100. Using AI and machines to perform tasks that were previously done by humans, potentially replacing some human jobs while creating new ones that didn't exist before
Matching (AI Ethics & Bias): 1→C, 2→D, 3→E, 4→A, 5→B
101. C — Explain that movie AI is science fiction — real AI can only do specific tasks it's trained for and doesn't have goals, desires, or consciousness to 'take over' anything
102. B — The measurable characteristics of the data that the AI uses to make predictions — like color, size, weight, or pixel values
103. D — When certain groups, perspectives, or conditions are underrepresented or missing from training data, causing the AI to perform poorly for those groups
104. B — Classification labels the whole image with one category, while object detection finds and locates multiple objects within an image, drawing bounding boxes around each one
105. C — A systematic evaluation of an AI system's outputs across different groups to check for bias, errors, and fairness — done regularly because bias can emerge

over time as data and society change

106. A — Any information that can identify you or be linked back to you as a specific person
107. Someone who uses technology responsibly, respects others' rights and privacy online, thinks critically about digital information, and contributes positively to online communities
108. Carefully crafted modifications to images that are imperceptible to humans but cause AI models to make confident but completely wrong classifications
109. Creating a shorter version of a text that captures the key points. Extractive summarization selects and combines the most important existing sentences; abstractive summarization generates entirely new sentences that convey the meaning
110. AI prompt engineer, data labeler/annotator, AI ethics officer, AI trainer/evaluator, machine learning engineer, AI safety researcher, conversational AI designer, and AI content moderator
Matching (Natural Language Processing): 1→A, 2→E, 3→B, 4→D, 5→C
111. B — First: the term 'Artificial Intelligence' was coined (1956); Then: chess AI beat a world champion (1997); Then: voice assistants like Siri launched (2011); Then: ChatGPT made AI widely accessible (2022)
112. B — The ability of AI to identify objects, people, or patterns in images — it works by analyzing thousands of pixel patterns and comparing them to patterns learned from training images
113. D — Verify AI outputs before sharing them as facts, cite AI when using its help, understand its limitations, consider who might be affected by AI decisions, and advocate for fair AI systems
114. B — A feature that correlates with a protected characteristic (like race or gender) even though the AI doesn't use that characteristic directly — zip codes, names, and purchasing patterns can serve as proxies
115. C — The ability to understand how an AI reached a specific decision — important because people affected by AI decisions deserve to know why, and hidden biases can only be fixed if they're visible
116. A — It requires two different types of verification (like a password plus a code sent to your phone), so even if someone steals your password, they still can't access your account
117. Privacy is the right to control what personal information you share and with

whom; secrecy is hiding something that might be harmful or deceptive

118. Using a model already trained on a large general dataset (like ImageNet) as a starting point, then fine-tuning it for a specific task with a much smaller dataset — saving enormous time and computing resources
119. An AI model trained on billions of text examples that generates responses by predicting the most likely next word in a sequence, based on patterns learned from its training data
120. For: AI can process more data than humans and reduce individual human bias.
Against: training data reflects historical discrimination, AI can't account for individual circumstances, and being labeled 'high risk' by an AI can become a self-fulfilling prophecy
Matching (Computer Vision Fundamentals): 1→C, 2→D, 3→B, 4→E, 5→A
121. A — Accuracy is the percentage of correct predictions — but it can be misleading when classes are imbalanced (e.g., 99% accuracy on disease detection when 99% of people are healthy just means it always says 'healthy')
122. D — AI systems analyze your behavior to build a profile of your interests and preferences, then use that profile to decide what content, ads, and recommendations to show you — creating a personalized but potentially limited view
123. D — A mathematical technique that adds carefully calibrated random noise to data or query results, making it impossible to identify any specific individual while preserving accurate statistical patterns about the overall population
124. B — Equality means treating everyone the same regardless of circumstances; equity means giving different people what they need to achieve fair outcomes — which sometimes requires treating them differently
125. D — AI can describe scenes, read text aloud, identify objects and people, navigate obstacles, and interpret facial expressions — acting as a 'visual assistant' that translates the visual world into audio or haptic information
126. B — When AI confidently generates false, fabricated, or misleading information that sounds plausible — this happens because the model predicts statistically likely text rather than verified facts
127. A government program giving every citizen a regular cash payment regardless of employment status — proponents argue it could cushion the economic disruption as AI automates jobs faster than new ones are created, ensuring everyone can meet basic needs during the transition
128. A standardized document accompanying an AI model that describes its intended

use, performance metrics across different demographics, limitations, training data information, and ethical considerations — similar to a nutrition label for food

129. AI built for one purpose (intended use) is often repurposed for different, potentially harmful purposes (actual use) — facial recognition built for unlocking phones repurposed for mass surveillance, for example
130. A strong project would identify a real problem (like identifying plant diseases in a school garden), specify what training data is needed (photos of healthy vs. diseased plants), acknowledge potential biases (seasonal variations, lighting differences), and define clear success metrics (accuracy rate, false positive rate, comparison with expert identification)
Matching (AI Basics): 1→A, 2→B, 3→E, 4→C, 5→D
131. D — AI learns by processing massive amounts of data to find statistical patterns; humans learn through experience, explanation, reasoning, and can learn from very few examples
132. B — When an AI model memorizes the training data too precisely instead of learning general patterns — it performs perfectly on training data but poorly on new, unseen data
133. B — AI-generated fake images, audio, or video that are realistic enough to fool people — raising concerns about misinformation, identity theft, harassment, and erosion of trust in authentic media
134. B — Draw many examples of each shape on cards with labels, show them to your friend one at a time as training data, then test with new unlabeled drawings to see if they've learned the patterns — this mimics supervised learning
135. D — Training large AI models consumes enormous energy — a single large language model training run can emit as much carbon as five cars over their lifetimes — raising questions about sustainability and who bears the environmental burden
136. D — QA systems understand the meaning of a question and extract or generate a specific answer from a knowledge source — unlike keyword search which just finds documents containing matching words without understanding what's being asked
137. A design approach where humans review, approve, or override AI decisions before they're acted upon — it's essential for high-stakes decisions where errors have serious consequences and where context, empathy, or ethical judgment are required

138. A design methodology that proactively integrates human values (fairness, privacy, autonomy, trust) into the technology design process from the very beginning — unlike 'ethics bolted on' at the end, it shapes fundamental design decisions and trade-offs throughout development
139. Understanding how AI works, its capabilities and limitations, how it affects your life, and being able to participate meaningfully in societal decisions about AI governance — this is essential because AI increasingly shapes what you see, what opportunities you get, and how you're treated by institutions
140. Your phone uses GPS, Wi-Fi networks, and cell towers to pinpoint your location, creating a detailed record of everywhere you go that apps and companies can access
Matching (AI Basics): 1→A, 2→C, 3→E, 4→D, 5→B
141. B — Large amounts of data (examples) and an algorithm (set of rules) that finds patterns in that data
142. B — When a biased AI's decisions create data that reinforces the original bias — for example, biased policing AI sends more officers to certain neighborhoods, creating more arrests there, which the AI uses as evidence to send even more officers
143. A — To check if the AI actually learned general patterns (not just memorized the training examples) — testing data the AI has never seen reveals its real-world performance
144. C — The task of assigning a label or category to an entire image, such as identifying whether a photo contains a cat, dog, or bird
145. A — Often poorly — idioms like 'it's raining cats and dogs' might be translated literally word-by-word instead of conveying the intended meaning, because idioms are culturally specific and their meaning can't be derived from individual words
146. C — AI can improve diagnosis accuracy, accelerate drug discovery, personalize treatment plans, and expand access to medical expertise — but risks include algorithmic bias in medical decisions, over-reliance on AI diagnosis, privacy concerns with health data, and potential to widen healthcare inequality
147. Metadata is 'data about data' — like the time, location, and recipient of a message rather than its content — and patterns in metadata can reveal detailed information about your life even without seeing the actual content
148. If a self-driving car must choose between two harmful outcomes (like swerving to avoid pedestrians but hitting a cyclist), who programs those life-or-death

decisions and by what values?

149. A U.S. law that protects children under 13 from having their personal data collected online without parental consent
150. The ability to detect and track the positions of human body joints (shoulders, elbows, knees, etc.) in images or video, mapping them as a skeleton
Matching (AI Basics): 1→D, 2→B, 3→A, 4→C, 5→E
151. B — Different lighting conditions, angles, juvenile vs adult appearance, similar-looking species, obstructed views, and ensuring the training data includes all species you want to classify
152. D — Privacy is a fundamental right that protects everyone, not just people with secrets; it prevents abuse of power, protects free expression, and data collected today could be misused by future governments or companies in ways we can't predict
153. B — Using knowledge from a model trained on one task to help with a different but related task — saving time and data because the model doesn't have to learn everything from scratch
154. B — Language is ambiguous (words have multiple meanings), context-dependent, culturally influenced, and constantly evolving — unlike chess which has fixed rules, a finite board, and clearly defined winning conditions
155. D — AI enables precision agriculture (targeted water/fertilizer use), crop disease detection from drone imagery, yield prediction, automated harvesting, and livestock health monitoring — potentially increasing food production while reducing environmental impact, but with risks of displacing small farmers and increasing dependence on technology
156. A — Most share core principles of: transparency, fairness and non-discrimination, accountability, privacy and data protection, safety and robustness, human oversight, and beneficence (AI should benefit humanity)
157. A small number of companies (Google, Microsoft, Meta, Amazon, Apple) control most AI research, talent, computing resources, and data — this concentration could lead to monopolistic behavior, reduced innovation, value systems of a few shaping technology used by billions, and insufficient democratic oversight
158. The field of AI focused on enabling computers to understand, interpret, and generate human language in a useful way
159. AI enables personalized learning at each student's pace, automated feedback on assignments, intelligent tutoring systems, and content generation for teachers — while raising questions about academic integrity, screen time, critical thinking

development, and the irreplaceable value of human teachers

160. It scrambles your data into unreadable code that can only be decoded with the correct key

Matching (AI Basics): $1 \rightarrow A$, $2 \rightarrow C$, $3 \rightarrow B$, $4 \rightarrow E$, $5 \rightarrow D$

161. To improve its learning rate!

162. A path-o-logical thinker!

163. It needed deeper connections!

164. Algo-rhythm!

165. It was trained on a lot of input!

166. "You only see what I show you!"

Spark & Anvil is a 501(c)(3) public charity. Apps + books are free forever.