

AI-900 Master Cheat Sheet

Describe Artificial Intelligence workloads and considerations (15–20%)

Identify features of common AI workloads

Content Moderation and Personalization Workloads:

- Features:
 - Detect and filter harmful or inappropriate content (images, text, video).
 - Tailor content and recommendations to individual users based on their preferences and behavior.
 - Common services: Azure Content Moderator, Azure Personalizer.

Computer Vision Workloads:

- Features:
 - Analyze and extract information from visual content (images, videos).
 - Identify objects, faces, text, and scenes.
 - Track objects and motion.
 - Common services: Azure Computer Vision, Azure Video Analyzer.

Natural Language Processing (NLP) Workloads:

- Features:
 - Understand and process human language.
 - Extract meaning, sentiment, entities, and relationships from text.
 - Translate languages.
 - Generate text summaries, question-answering, and chatbots.
 - Common services: Azure Text Analytics, Azure Translator, Azure Language Understanding (LUIS).

Knowledge Mining Workloads:

- Features:
 - Extract knowledge from unstructured text sources (documents, emails, web pages).
 - Discover patterns, relationships, and insights.
 - Build knowledge graphs and ontologies.
 - Common services: Azure Cognitive Search, Azure Machine Learning.

Document Intelligence Workloads:

- Features:
 - Analyze and process structured and semi-structured documents (invoices, forms, contracts).
 - Extract key information (text, tables, fields).
 - Automate document processing tasks.
 - Common services: Azure Form Recognizer, Azure Document Translation.

Generative AI Workloads:

- Features:
 - Create new content (text, images, audio, video) based on patterns learned from existing data.
 - Generate realistic and creative outputs.
 - Common services: Azure OpenAI Service (powered by OpenAI's models).

Identify guiding principles for responsible AI

Fairness:

- **Bias:** Identify and mitigate potential biases in training data and algorithms. Ensure the solution doesn't unfairly advantage or disadvantage any particular group.
- **Transparency:** Explain how decisions are made and why certain outcomes are reached. This allows for identification and correction of unfairness.
- **Accessibility:** Ensure everyone can access and benefit from the AI solution. Consider language, education level, and physical abilities.

Reliability and Safety:

- **Data Quality:** Ensure data used for training and testing is accurate, complete, and relevant. Poor data leads to unreliable results.
- **Robustness:** Test the solution for errors and unexpected input. Make sure it can handle failures gracefully and minimize risks.
- **Monitoring:** Continuously monitor the AI solution for performance and potential negative impacts. Have protocols in place to address issues quickly.

Privacy and Security:

- **Data Protection:** Implement robust security measures to protect user data from unauthorized access, theft, or misuse. Ensure compliance with data privacy regulations.
- **Transparency:** Inform users about what data is collected, how it is used, and their rights concerning their data.

- **Control:** Give users control over their data and how it is used in the AI solution. Allow them to opt-in/out and access/correct their data.

Inclusiveness:

- **Accessibility:** Ensure the solution is accessible to people with diverse abilities and backgrounds. Consider language, cultural differences, and accessibility needs.
- **Representation:** Include diverse perspectives in the design and development process to avoid reinforcing existing biases or excluding certain groups.
- **Benefit:** Ensure the AI solution benefits everyone equally and does not disproportionately harm any particular group.

Transparency:

- **Explainability:** Explain how the AI solution works and reaches its decisions. This helps users understand the reasoning behind the outputs and builds trust.
- **Algorithm Transparency:** If possible, reveal the algorithms used in the solution, or at least offer high-level explanations of their function.
- **Communication:** Clearly communicate the capabilities and limitations of the AI solution to users. Avoid overhyping or making misleading claims.

Accountability:

- **Attribution:** Clearly identify who is responsible for the development, deployment, and use of the AI solution. This ensures accountability for potential harms.
- **Auditing:** Regularly audit the AI solution for ethical lapses and compliance with regulations. Have mechanisms for reporting and addressing concerns.
- **Remediation:** Have clear protocols for addressing harmful outcomes or unintended consequences of the AI solution. This includes providing compensation or fixing errors.

Describe fundamental principles of machine learning on Azure (20–25%)

Identify common machine learning techniques

Machine Learning Scenarios

Machine learning involves training algorithms to identify patterns and make predictions from data. Choosing the right type of algorithm depends on the kind of problem you're trying to solve. Here's a breakdown of different scenarios and their corresponding machine learning models:

Regression Machine Learning:

- **Scenario:** You want to predict a continuous value based on other features.
- **Examples:** Forecasting future temperatures, estimating house prices, predicting stock market trends.

- Models: Linear regression, polynomial regression, support vector regression, random forest regression.

Classification Machine Learning:

- Scenario: You want to predict a discrete category based on other features.
- Examples: Classifying emails as spam/not spam, identifying fraudulent transactions, recognizing handwritten digits.
- Models: Logistic regression, K-nearest neighbors, decision trees, support vector machines, random forests.

Clustering Machine Learning:

- Scenario: You want to group unlabeled data points based on their similarities.
- Examples: Identifying customer segments, finding anomalies in sensor data, grouping genes with similar functions.
- Models: K-means clustering, hierarchical clustering, density-based spatial clustering (DBSCAN).

Features of Deep Learning Techniques

Deep learning is a specific type of machine learning that uses artificial neural networks inspired by the structure and function of the brain. Some key features include:

- Non-linearity: Deep learning models can learn complex, non-linear relationships between features, unlike simpler models like linear regression.
- High feature extraction: Deep neural networks can automatically extract relevant features from raw data, reducing the need for manual feature engineering.
- High dimensionality: They can handle data with a large number of features effectively.
- Representation learning: They learn internal representations of the data that are more informative than the raw features.
- Scalability: They can be trained on large datasets, leading to better performance with more data.

In summary:

- Choose regression for predicting continuous values, classification for discrete categories, and clustering for grouping unlabeled data.
- Deep learning offers advantages like non-linearity, automatic feature extraction, and high dimensionality handling, but requires more data and computational resources.

Describe core machine learning concepts

Features and Labels:

- Features are the individual input variables that represent the characteristics or attributes of the data points in a dataset. They are used by machine learning models to make predictions or decisions.
- Labels, also known as targets or ground truth, are the known outcomes or values that correspond to each data point. They are used to train supervised machine learning models, which learn to map features to labels.

Example:

Consider a dataset of images of cats and dogs. The features might include pixel values, color information, texture details, etc. The labels would be "cat" or "dog," indicating the correct classification for each image.

Training and Validation Datasets:

- Training dataset is a portion of the overall dataset that is used to train a machine learning model. The model learns to identify patterns and relationships between features and labels within the training data.
- Validation dataset is a separate portion of the dataset that is not used for training. It is used to evaluate the model's performance on unseen data and to fine-tune its hyperparameters. By evaluating model performance on the validation set, we can ensure that the model is not overfitting to the training data.

Describe Azure Machine Learning capabilities

1. Automated Machine Learning (AutoML):

AutoML streamlines the machine learning process by automating various tasks, making it accessible even for non-experts. Here's what it can do:

- Data preparation: Cleans, transforms, and features engineering your data automatically.
- Model training: Selects and trains various models based on your data and task, saving you time and effort.
- Model selection: Compares and ranks trained models based on their performance, helping you choose the best one.
- Hyperparameter tuning: Optimizes model hyperparameters for better accuracy and efficiency.
- Explainability: Generates insights into why a model makes certain predictions, improving transparency.

2. Data and Compute Services:

Azure offers a range of services for various data science and machine learning needs:

- Data storage: Store data in different formats (structured, unstructured) with services like Azure Blob Storage, Azure Data Lake Storage, and Azure SQL Database.

- **Data management:** Azure Synapse Analytics and Azure Data Factory provide tools for data ingestion, cleansing, transformation, and orchestration.
- **Compute:** Access powerful VMs, container instances, and serverless options like Azure Machine Learning compute or Azure Functions for training and deploying models.

3. Model Management and Deployment in Azure Machine Learning:

Azure Machine Learning simplifies the model lifecycle management:

- **Model registration:** Store and track different versions of your trained models.
- **Model packaging:** Prepare models for deployment in various environments (cloud, on-premises, edge devices).
- **Model deployment:** Deploy models to Azure web services, containers, or other platforms for real-world use.
- **Monitoring and evaluation:** Track model performance metrics and ensure continuous improvement.
- **Versioning and rollback:** Easily switch between model versions and roll back if needed.

Benefits of using these capabilities:

- **Faster time to insights:** Streamlines data handling and model training, accelerating your journey to actionable results.
- **Lower costs:** Reduces the need for specialized expertise and avoids expensive infrastructure overhead.
- **Improved model performance:** AutoML algorithms often outperform manually built models.
- **Increased productivity:** Data professionals can focus on higher-level tasks instead of routine chores.
- **Scalability and flexibility:** Azure's cloud services adapt to your growing needs and diverse deployment options.

By leveraging these capabilities, you can build and deploy efficient, data-driven solutions across various industries and applications.

Describe features of computer vision workloads on Azure (15–20%)

Identify common types of computer vision solution:

1. Image Classification:

- **Model performance:** Measures how accurately the model classifies images into predefined categories. Relevant metrics include accuracy, precision, recall, and F1 score.
- **Speed and latency:** Refers to how quickly the model processes and classifies an image. This is crucial for real-time applications.

- Scalability and adaptability: The ability to handle large datasets and adapt to new categories or changes in existing ones.
- Flexibility and customization: Whether the model can be tailored to specific requirements and datasets for improved performance.
- Interpretability: The ability to understand the model's reasoning behind its classification, enhancing trust and making debugging easier.

2. Object Detection:

- Multi-object detection: The ability to identify and localize multiple objects within an image, not just one.
- Bounding box accuracy: How precisely the model can delineate the location and size of each detected object.
- Background suppression: Effectively differentiating objects from the background for accurate detection.
- Class-specific features: Recognizing specific features and attributes of objects based on their category.
- Real-time tracking: Ability to track identified objects continuously across video frames.

3. Optical Character Recognition (OCR):

- Accuracy and precision: Recognition of characters (including special symbols) with minimal errors.
- Font and style recognition: Adaptability to different font types, sizes, and styles.
- Language support: The ability to recognize characters in multiple languages.
- Complex document handling: Recognition of text within images of documents with layouts, tables, and other formatting elements.
- Integration with other systems: Seamless integration with OCR results into downstream workflows.

4. Facial Detection and Facial Analysis:

- Robustness: Accurate detection of faces under various lighting conditions, poses, and occlusions.
- Multi-face detection: Ability to identify and locate multiple faces within an image.
- Facial landmark detection: Recognizing key points on the face (eyes, nose, mouth) for further analysis.
- Emotion recognition: Ability to infer emotions from facial expressions.
- Biometric identification: Matching faces to known individuals for identification or verification.

Identify Azure tools and services for computer vision tasks

1. Azure AI Vision:

- **Image Analysis:** Identifies objects, scenes, and actions in images, along with emotions and facial expressions. Extracts text from images using optical character recognition (OCR).
- **Image Classification:** Categorizes images into pre-defined categories or custom categories you create.
- **Tagging:** Generates descriptive tags for images based on their content.
- **Custom Vision:** Train your own custom image classification models based on your specific needs.
- **Object Detection:** Locates and identifies objects within images and provides bounding boxes around them.

2. Azure AI Face Detection:

- **Face Detection:** Locates faces in images and videos with high accuracy, even in challenging conditions like occlusion or low resolution.
- **Facial Recognition:** Identifies known individuals in images and videos based on a pre-trained face database.
- **Facial Analysis:** Extracts emotional attributes, age, gender, and other facial features from detected faces.
- **Person Identification:** Tracks identified individuals across multiple images and videos.
- **Large Gallery Search:** Identifies known individuals in a large gallery of faces with millions of entries.

3. Azure AI Video Indexer:

- **Automatic Speech Recognition (ASR):** Transcribes spoken audio in videos to text with support for multiple languages and accents.
- **Optical Character Recognition (OCR):** Extracts text from images within videos, such as captions or signs.
- **Scene Detection:** Automatically identifies and segments different scenes within a video based on changes in content or camera angles.
- **Topic Modeling:** Analyzes video content and identifies key topics discussed.
- **Sentiment Analysis:** Understands the sentiment (positive, negative, or neutral) of spoken dialogue and overall video content.
- **Speaker Identification:** Recognizes and distinguishes different speakers within a video based on their voices.
- **Content Moderation:** Detects potentially inappropriate content such as nudity, violence, or profanity in videos.

Describe features of Natural Language Processing (NLP) workloads on Azure (15–20%)

Identify features of common NLP Workload Scenarios

1. Key Phrase Extraction:

- Features:
 - Identifies the most important and relevant concepts, keywords, and topics within a text.
 - Uses statistical methods, linguistic rules, and machine learning algorithms for analysis.
 - Can be customized to target specific domains or tasks.
- Uses:
 - Summarization of documents and news articles.
 - Information extraction and retrieval from text.
 - Clustering and tagging documents based on content.
 - Keyword research for marketing and SEO.

2. Entity Recognition:

- Features:
 - Identifies and classifies named entities like people, organizations, locations, dates, and quantities within text.
 - Employs gazetteers, dictionaries, and machine learning models for classification.
 - Can be extended to identify custom entity types relevant to specific domains.
- Uses:
 - Question answering systems and conversational bots.
 - Information extraction and knowledge base construction.
 - Financial analysis and risk assessment.
 - Sentiment analysis based on entity context.

3. Sentiment Analysis:

- Features:
 - Analyzes the emotional tone and opinion expressed in text data.
 - Classifies sentiment as positive, negative, neutral, or mixed.
 - Utilizes lexical analysis, machine learning, and rule-based approaches.

- Can be fine-tuned to analyze specific emotions like anger, joy, or sadness.
- Uses:
 - Customer feedback analysis and social media monitoring.
 - Market research and product development.
 - Political opinion analysis and risk assessment.
 - Personalized recommendations and targeted advertising.

4. Language Modeling:

- Features:
 - Predicts the next word in a sequence based on the preceding context.
 - Enables generation of coherent and grammatically correct text.
 - Employs statistical models like n-grams or neural networks for prediction.
 - Can be trained on various text formats and genres.
- Uses:
 - Machine translation and text summarization.
 - Dialogue systems and chatbot development.
 - Automated writing and content creation.
 - Text prediction and autocompletion in software applications.

5. Speech Recognition and Synthesis:

- Speech Recognition:
 - Features:
 - Converts spoken language into text format.
 - Utilizes acoustic models and language models for accurate transcription.
 - Adapts to different accents and speaking styles.
 - Uses:
 - Voice assistants and smart home devices.
 - Dictation and voice-to-text applications.
 - Automated call centers and speech analytics.
- Speech Synthesis:
 - Features:
 - Converts text into spoken language.
 - Generates natural-sounding audio with different voices and emotions.

- Uses:
 - Text-to-speech applications for accessibility and convenience.
 - Interactive audiobooks and voiceovers.
 - Chatbots and virtual assistants with more engaging interactions.

6. Translation:

- Features:
 - Converts text from one language to another while preserving meaning.
 - Utilizes statistical machine translation and neural networks for accuracy.
 - Adapts to different domains and contexts.
- Uses:
 - Website and document localization for global audiences.
 - Real-time communication and interpretation tools.
 - Content creation and dissemination across languages.
 - Breaking down language barriers in information access and communication.

Identify Azure tools and services for NLP workloads

1. Azure AI Language Service:

- Text Analytics: Analyze text for sentiment, key phrases, entities, language, and more. Extract information and gain insights from unstructured text data.
- Text Mining: Uncover hidden relationships and patterns within large volumes of text documents. Perform topic modeling, clustering, and anomaly detection.
- Question Answering: Build intelligent chatbots and virtual assistants that can answer questions about your data in a natural way.
- Optical Character Recognition (OCR): Extract text from images and documents with high accuracy, supporting various languages and fonts.
- Custom Vision: Train your own image classification models with custom datasets to recognize specific objects, scenes, or actions in images.
- Machine Translation: Translate text between languages with high-quality output, supporting dozens of languages and dialects.

2. Azure AI Speech Service:

- Speech-to-Text: Convert spoken audio into text with real-time or batch processing capabilities. Ideal for transcribing audio recordings, podcasts, and conversations.

- **Text-to-Speech:** Synthesize realistic human-like speech from text in various languages, accents, and voices. Generate audio content for your applications, videos, or e-learning modules.
- **Speaker Recognition:** Identify individual speakers in audio recordings based on their unique vocal characteristics. Useful for customer service applications, meeting transcripts, and security systems.
- **Sentiment Analysis:** Analyze the emotional tone and sentiment of spoken language in real-time, providing insights into customer interactions or call center conversations.
- **Custom Speech:** Train your own speech recognition models to recognize specialized vocabulary or domain-specific language relevant to your industry or application.

3. Azure AI Translator Service:

- **Machine Translation:** Translate text and speech between over 70 languages and dialects with high accuracy and fluency. Integrate seamless real-time or batch translation into your applications.
- **Custom Translator:** Train your own translation models based on your specific terminology and domain knowledge for even more accurate and nuanced translations.
- **Speech Translation:** Translate live conversations in real-time between multiple languages, enabling multilingual communication and collaboration.
- **Document Translation:** Translate entire documents, including PDFs, Word files, and web pages, while preserving formatting and layout.
- **Neural Machine Translation (NMT):** Utilize the latest NMT technology for state-of-the-art translation quality, delivering natural-sounding and accurate results.

Describe features of generative AI workloads on Azure (15–20%)

Identify features of generative AI solutions

Generative AI: Features, Scenarios, and Responsible Considerations

Generative AI is a branch of artificial intelligence focused on creating new content, data, or code. Unlike traditional AI models that analyze or classify existing information, generative models learn from existing data and use that knowledge to produce something entirely new.

Features of Generative AI Models:

- **Creativity:** They can generate novel content that wasn't explicitly taught in the training data, mimicking human creativity in areas like text, music, art, and code.
- **Variability:** The same model can produce different outputs from the same input, offering diverse and original results.
- **Scalability:** They can generate large amounts of content quickly and efficiently, making them ideal for tasks like data augmentation or personalization.

- **Adaptability:** Can be readily fine-tuned and customized to specific applications or purposes.

Common Scenarios for Generative AI:

- **Text Generation:** Creating realistic dialogue for chatbots, writing compelling marketing copy, and generating creative fiction.
- **Image and Video Synthesis:** Creating realistic images and videos from text descriptions, editing existing visuals, and even generating deepfakes.
- **Music Composition:** Generating new music in different styles, producing sound effects, and even remixing existing songs.
- **Software and Code Generation:** Automatically writing code for specific tasks, generating test data for software development, and even composing complete applications.

Responsible AI Considerations for Generative AI:

- **Bias and Fairness:** Generative models can inherit biases from the data they're trained on, leading to discriminatory outputs. Careful data selection and mitigation strategies are crucial.
- **Transparency and Explainability:** It's often difficult to understand how generative models reach their outputs, making it challenging to address issues like bias or ensure the accuracy of generated content.
- **Misinformation and Disinformation:** Malicious actors can use generative AI to create and spread false information, highlighting the need for robust detection and verification mechanisms.
- **Ownership and Copyright:** Questions arise about who owns the copyright of content generated by AI models, requiring clear legal frameworks and ethical considerations.

Identify capabilities of Azure OpenAI Service

Natural Language Generation (NLG) Capabilities:

- **Text Generation:** Create diverse text formats, including:
 - Emails, letters, blog posts, social media content
 - Scripts, poems, code, musical pieces
 - Product descriptions, marketing materials
 - Answers to questions, summaries of text
 - Creative text formats, like poems, code, scripts, musical pieces, email, letters, etc.
- **Conversational AI:** Build chatbots and virtual assistants that engage in natural, fluid conversations.
- **Translation:** Translate text between languages with fluency and accuracy.

- Content Creation: Generate content for various purposes, such as marketing, education, and entertainment.

Code Generation Capabilities:

- Generate code snippets: Create functions, methods, classes, or even complete programs in various programming languages.
- Translate natural language to code: Convert descriptions of desired code functionality into actual code.
- Write unit tests: Automate test generation for code quality assurance.
- Assist with code completion: Help developers by suggesting code completions and autocomplete options.
- Refactor code: Improve code readability and maintainability.

Image Generation Capabilities:

- Create images from text descriptions: Generate illustrations, photographs, and photorealistic images based on text prompts.
- Edit and manipulate existing images: Change image features, style, perspective, or add/remove elements.
- Combine text and images: Generate images that incorporate text elements, such as captions, logos, or text-based art.
- Generate images for various purposes: Create product mockups, social media visuals, design prototypes, or artistic imagery.

Disclaimer: All data and information provided on this site is for informational purposes only. This site makes no representations as to accuracy, completeness, correctness, suitability, or validity of any information on this site & will not be liable for any errors, omissions, or delays in this information or any losses, injuries, or damages arising from its display or use. All information is provided on an as-is basis.