AI Use Case Ideation Canvas

This canvas is designed to guide you through a structured process to ideate, conceptualize, and define AI use cases that are innovative, user-centric, feasible, and aligned with organizational goals. By providing a comprehensive overview of AI capabilities and prompting thoughtful consideration of user needs, organizational processes, available data, and technological feasibility, this canvas ensures that the ideated AI use cases are grounded in real-world problems and opportunities and hold the promise of delivering tangible value and impact. Utilizing this canvas is crucial for fostering a deep understanding of AI's potential and limitations and for enabling informed decision-making during the subsequent evaluation and implementation of the AI use cases.

Name:

Team Name:

Date:

AI Use Case Canvas

Use Case Name

Use Case Description

Provide a brief summary of the use case, highlighting the purpose and key outcomes of the AI solution.

- What problem does this AI solution solve?
- How does it improve a current process or user experience?

Problem

Clearly describe the problem or opportunity this AI use case addresses. This should include specific inefficiencies or gaps that the solution aims to solve.

- What are the current pain points or inefficiencies?
- How does this problem affect productivity, decision-making, or outcomes?

Input

Specify the data, information, or signals that the AI solution requires to function. This could include text, images, audio, structured data, etc.

- What data or inputs are needed to power this AI solution?
- Where does the data come from (e.g., sensors, databases, human input)?

Output

Define the results or outcomes the AI solution produces. This could be a decision, prediction, report, or other actionable output.

- What is the main output or result the AI solution will produce?
- How is this output delivered or made available (e.g., report, notification, dashboard)?

Area of Application

Define whether the AI use case applies to a product, service, or process, and specify the context.
Product/Service: Indicate whether this is part of enhancing an existing product/service or a new offering.
Process: Specify if it's a general, industry-specific, function-specific, or unique process.

- Is this AI solution part of a product, service, or process? If a process, is it specific to a function, industry, or general?
- Does this solution create something new or enhance an existing offering?

AI Capability Used

List the AI capabilities that will be leveraged to make this use case possible. Include key technologies or methods like computer vision, natural language processing, or machine learning.

- Which AI capabilities are required to implement this solution?
- How will these technologies solve the problem or improve the process?

Benefits

Estimate the benefit of the use case by evaluating its strategic impact (alignment with business goals) and economic value (cost savings, revenue potential). Rank each aspect from 0-5 and give it an overall score.

- **Strategic Benefit (0-5):**
  How well does the use case align with long-term organizational strategy and innovation?
- **Economic Benefit (0-5):**
  What is the potential for cost savings or revenue generation?
- **Overall Score:**
  A weighted score considering both strategic and economic benefits.

Cost

Rank the complexity and cost of implementing the AI use case. A high score indicates high cost or difficulty. Rank both on a scale from 0-5, and calculate a weighted overall score.

- **Implementation Cost (0-5):**
  What is the estimated cost of building, maintaining, and scaling this AI solution?
- **Ease of Implementation (0-5):**
  How complex is the technical development and integration into existing systems?
- **Overall Score:**
  A weighted score considering both cost and ease of implementation

Potential Risks

Identify risks that could complicate the implementation or lead to reputational or legal challenges. Include regulatory, ethical, technical, and operational risks.
Guiding Questions:

- Could this AI use case introduce reputational, privacy, or ethical risks?
- What technical or regulatory challenges might hinder the successful deployment of this AI solution?

Triggered Action

Outline the actions or decisions that the AI solution will enable or automate once it generates the output. This could involve a process being triggered, an alert being sent, or an action being recommended or taken.

- What actions or processes will be triggered by the AI solution's output?
- Will this output trigger an automated or manual response?

Value Generated

Describe the tangible or intangible value that the AI solution will create for the organization or users. This could include time savings, improved decision-making, or enhanced efficiency.

- What is the primary value this AI solution brings to the organization or users?
- How does the value align with strategic goals, and can it be measured?

Key Hypotheses to Validate

Outline the key hypotheses that need to be tested before fully committing to the AI use case. These should focus on assumptions related to the solution's effectiveness, feasibility, and potential impact.

- What assumptions about this AI use case must be validated to ensure its success?
- What would be the key indicators that this use case is solving the intended problem effectively?

Market Solutions (Make vs. Buy)

Research existing AI solutions available in the market. List 1-2 solutions and assess whether they meet the use case's requirements. This helps determine if the solution should be developed internally (make) or purchased externally (buy).

- Are there existing solutions in the market that fulfill this use case?
- If yes, are they scalable and adaptable to your organization's needs? Or is an in-house solution necessary? Available Solutions: Indicate the name(s) of relevant solutions and whether they should be further assessed or if this is a case where no suitable solution exists yet.

Overview of AI Capabilities:

AI Capability:
A specific function or set of functions that an artificial intelligence system can perform, achieved through learning from data and utilizing computational algorithms, allowing the system to execute tasks traditionally requiring human intelligence .

Perception Capabilities:
These capabilities enable machines to interpret and understand the world around them. They involve sensory data input, such as visual, auditory, or linguistic, and are akin to human senses. Examples include computer vision, audition, and linguistics.

Analytical Capabilities:
Analytical capabilities empower AI to process, evaluate, and draw conclusions from data. These functionalities are centered around finding patterns, predicting future events, and making optimal decisions. Key areas are discovery, forecasting, and planning & optimization.

Motoric Capabilities:
Motoric capabilities provide AI systems the ability to interact physically with their environment. These encompass movement, manipulation, and dynamic response. The core representation of this category is advanced robotics.

Generating Capabilities:
These capabilities allow AI to create new content or simulations based on existing data or patterns. They're focused on producing unique outputs, whether that's artwork, music, or novel data structures. Generative AI is a prime example.

Computer Vision:

Computer vision involves teaching machines to interpret and act based on visual data. By analyzing images and videos, these systems can recognize objects, track movements, and even gauge emotions. It's the technology behind facial recognition, autonomous vehicles, and augmented reality.

Object Detection and Tracking:

Description: Object detection is about identifying specific items within an image or video, while tracking follows their movement over time. Advanced algorithms can discern between multiple objects and types within a scene.

Example: Security cameras in retail stores use object detection to identify potential shoplifters, and then track their movements throughout the store to gather evidence or alert security personnel.

Image Classification:

Description: Image classification assigns a label to an entire image or photograph based on its overall content. It categorizes visuals into predefined classes.

Example: Social media platforms, like Instagram, might use image classification to automatically tag photos as 'beach', 'forest', or 'city' based on their content, aiding in search and content discovery.

Emotion Recognition:

Description: By analyzing facial features, emotion recognition systems can infer a person's emotional state from an image or video feed.

Example: Companies may use emotion recognition during product testing or ad screenings. When participants watch a new advertisement, a camera captures their facial reactions, helping marketers understand which parts of the ad evoke joy, surprise, or sadness.

3D Reconstruction:

Description: 3D reconstruction involves capturing the shape and appearance of real-world objects or environments and transforming this data into a three-dimensional digital model.

Example: In the real estate industry, agents might use 3D reconstruction to create virtual tours of properties. Cameras capture every angle of a home, and software creates a digital model that online users can "walk through" without visiting in person.

Image Segmentation:

Description: Image segmentation partitions an image into multiple segments or sets, making objects or boundaries stand out from the background and each other.

Example: Medical imaging often employs image segmentation. For instance, in an MRI of the brain, segmentation can separate the brain's different regions or highlight tumors, making them distinct from healthy tissue.

Computer Audition:

Computer audition allows machines to process, analyze, and understand audio signals. It involves recognizing speech, identifying sounds, and processing music. Applications include voice assistants, sound anomaly detection in industries, and music recommendation systems.

Speech to Text:

Description: Speech to text, also known as automatic speech recognition (ASR), converts spoken language into written text. It's powered by algorithms that learn from vast amounts of spoken data and their corresponding transcriptions.

Example: Voice assistants like Amazon's Alexa or Apple's Siri utilize speech-to-text to transcribe user commands. This allows users to set reminders, send messages, or search the internet just by speaking.

Musical Knowledge:

Description: Musical knowledge in AI encompasses the understanding and recognition of musical elements, such as pitch, tempo, and genre. Systems can be trained to identify songs, categorize music, or even create new compositions.

Example: Apps like Shazam use this capability to identify songs playing in the background. Users simply play a snippet of music, and the app can name the song and artist.

Sound Similarity Assessment:

Description: This involves comparing audio files or sounds to determine how similar they are. It can be used for matching, categorization, or even anomaly detection in soundscapes.

Example: Streaming services like Spotify might employ sound similarity assessment to suggest songs that have a similar audio "feel" to what a user has previously enjoyed, thereby curating personalized playlists.

Source Separation:

Description: Source separation aims to distinguish and isolate individual audio sources from a mixed audio recording. It's especially useful in noisy environments or where many sounds overlap.

Example: In a busy cafe, multiple conversations occur simultaneously. Using source separation, a smart hearing aid could isolate and amplify the voice of a person directly speaking to the hearing aid wearer, minimizing background chatter.

Audio-based Sentiment Analysis:

Description: It doesn't just transcribe what is said, but also how it's said, determining feelings like happiness, anger, or sadness.

Example: Call centers often use audio-based sentiment analysis. By monitoring customer calls, the system can detect when a customer sounds frustrated or angry, possibly alerting a supervisor or suggesting the agent switch tactics.

Computer Linguistics:

Also known as computational linguistics, this area focuses on the interaction between computers and human language. It enables machines to understand, generate, and respond to text or spoken words. It's foundational for chatbots, language translation, and sentiment analysis.

Translation:

Description: This involves converting text or speech from one language to another while maintaining its semantic meaning. Modern AI-powered systems can translate multiple languages in real-time.

Example: Google Translate is a popular tool that allows users to translate text, both written and spoken, into various languages. This tool is invaluable for travelers, students, and businesses operating internationally.

Text Classification:

Description: Text classification assigns predefined categories (or labels) to a given text based on its content. This can be used for spam detection, topic assignment, and more.

Example: : Email services like Gmail use text classification to filter and categorize incoming emails as "Primary", "Social", "Promotions", or even "Spam".

Sentiment Analysis:

Description: Sentiment analysis determines the emotional tone or subjective information behind a piece of text. It's commonly used to gauge public opinion or customer satisfaction.

Example: Companies often use sentiment analysis on product reviews to understand customer satisfaction levels. If a product consistently receives reviews mentioning it's "disappointing" or "frustrating", the sentiment would be negative.

Entity Recognition:

Description: Entity recognition identifies and categorizes specific entities within a text into predefined groups such as names of persons, organizations, locations, expressions of times, quantities, and more.

Example: In a news article about Apple launching a new product, entity recognition would identify "Apple" as an organization and the product name as a specific object or item.

Relation Extraction:

Description: This process aims to identify and extract relationships between named entities in the text. This can help in building knowledge graphs or understanding contextual associations.

Example: From the sentence "Barack Obama was born in Hawaii," relation extraction would identify the relationship between the entity "Barack Obama" and "Hawaii" as a birthplace association.

Conversational Systems:

Description: Conversational systems, or chatbots, simulate human conversation. They can understand user input, process it, and provide appropriate responses, enabling human-like interaction with software.

Example: Customer support chatbots are common on many websites. Users can ask questions or report issues, and the chatbot can assist in real-time, either solving the problem or directing the user to the right resource.

Discovery:

Discovery in AI refers to uncovering insights from vast amounts of data. By analyzing this data, AI can highlight hidden patterns, correlations, or anomalies. It's crucial for data mining, medical diagnoses, and fraud detection.

Segmentation and Clustering:

Description: This involves grouping data points based on their similarities without prior knowledge of categories. The goal is to identify inherent structures within data.

Example: SMarketing teams use clustering to segment customers into different groups based on their purchasing behaviors, enabling them to tailor marketing strategies to each group's preferences.

Anomaly/Outlier Detection:

Description: Anomaly detection identifies data points that deviate significantly from the expected pattern or the majority of data. It's essential for identifying rare events or potential errors.

Example: Credit card companies use anomaly detection to identify potential fraudulent activities. If someone usually makes purchases in Texas and suddenly there's a flurry of purchases in Paris, the system might flag it as suspicious.

Correlation Analysis:

Description: This process measures the degree to which two variables change in relation to each other. A strong correlation indicates that as one variable changes, the other is likely to follow a predictable pattern.

Example: In stock market analysis, analysts may study the correlation between a company's stock price and various economic indicators to predict future price movements.

Causal Inference:

Description: Causal inference determines the cause-and-effect relationship between variables. It goes beyond correlation to understand if one variable directly influences another.

Example: Medical researchers might use causal inference to determine if a specific drug reduces the risk of a disease. If those taking the drug have significantly fewer disease cases, and other factors are controlled for, there might be a causal relationship.

Association Analysis:

Description: Also known as market basket analysis, this technique identifies patterns of co-occurrence in datasets, often used to find items that are typically purchased together.

Example: Retail stores use association analysis to optimize product placement. If they discover that customers often buy chips and salsa together, they might place them close in the store or offer combined discounts.

Forecasting:

Forecasting uses historical data to predict future outcomes or trends. Through statistical models and algorithms, AI can anticipate stock market movements, weather patterns, or sales figures with increased accuracy.

Time Series Forecasting:

Description: Time series forecasting involves predicting future values based on past and present data points collected at successive time intervals. These predictions can be short-term or span across longer horizons, depending on the data and the application.

Example: Weather forecasting is a classic application of time series forecasting. Meteorologists analyze past and current weather data, such as temperature, humidity, and wind speed, to predict conditions for the coming days or weeks.

Dependency-based Forecasting:

Description: Dependency-based forecasting predicts one variable based on the values of one or more other variables. It identifies and leverages inter-variable relationships to generate forecasts, often employing regression models or similar techniques.

Example: In economics, dependency-based forecasting might be used to predict a country's future GDP growth based on factors like investment in infrastructure, education levels, and trade balances. If a country significantly invests in infrastructure, for instance, this could be correlated with a future rise in GDP.

Planning & Optimization:

This capability focuses on determining the best approach or solution for a given problem. AI can devise routes for delivery trucks, schedule tasks for optimal efficiency, or even strategize game moves.

Cooperative Multi-Agent Systems:

Description: This involves multiple independent agents that collaborate to achieve a shared objective. Each agent, while autonomous, is aware of the goals of the other agents and makes decisions that collectively benefit the group.

Example: In traffic management, multiple autonomous vehicles (agents) communicate with each other to optimize traffic flow. If one vehicle is slowing down due to an obstacle, it can inform other nearby vehicles, which can then adjust their routes or speeds to avoid congestion.

Policy Development/Strategic Agents:

Description: Strategic agents are designed to make decisions based on long-term goals and strategies. They analyze current states, consider various actions, and predict potential outcomes to devise and adapt policies.

Example: In energy management, strategic agents can determine when to store energy, when to use it, or when to sell it to the grid based on current energy prices, predicted future prices, and storage capacity.

Logistics Planning:

Description: Logistics planning deals with the coordination and optimization of resources and processes to transport goods or provide services. It encompasses route optimization, resource allocation, and inventory management.

Example: E-commerce giants like Amazon use logistics planning to determine the best routes for their delivery trucks, ensuring that packages are delivered to customers in the most efficient and timely manner.

Planning and Scheduling:

Description: This involves the arrangement, coordination, and timing of tasks and resources. The goal is to optimize certain objectives, such as minimizing costs or maximizing efficiency, given constraints.

Example: Airlines use planning and scheduling systems to assign pilots and crews to flights, ensuring they adhere to regulations regarding working hours and rest periods while maximizing aircraft utilization.

Advanced Robotics & Control:

Advanced robotics integrates AI with mechanical devices, enabling them to perform tasks autonomously in complex environments. From surgical robots to warehouse automation, these systems combine perception, decision-making, and precise motor control.

Robot Motion Planning:

Description: This involves creating paths for a robot to move from a starting position to a target position, avoiding obstacles and optimizing for certain criteria, such as shortest distance or energy efficiency.

Example: In automated manufacturing, robot arms use motion planning to pick up parts from one location and assemble them in another, ensuring they don't collide with other objects or machinery.

HD Mapping and Localization:

Description: High-Definition (HD) mapping provides detailed and accurate map data essential for applications like autonomous driving. Localization is the process of determining a robot's position within this map.

Example: Autonomous vehicles utilize HD maps that provide information beyond traditional maps, like lane specifics and road elevations. These vehicles constantly determine their position within this map, ensuring safe and accurate navigation.

Control Optimization:

Description: This involves designing control strategies that optimize certain performance criteria. It's about making dynamic systems, like robots or drones, behave in the best possible way under given conditions.

Example: Electric cars use control optimization to manage battery consumption, adjusting various parameters to ensure maximum driving range and battery health.

aborative Robotics/Human-Robot Interacti

Description: This focuses on robots designed to interact with humans in shared environments. They're built with safety features and sensors to recognize human gestures, voices, and movements.

Example: In some modern factories, collaborative robots work alongside human workers, assisting in tasks like heavy lifting or precision assembly, and adjusting their actions based on human co-worker movements.

Advanced Drones:

Description: These are unmanned aerial vehicles equipped with sophisticated sensors, cameras, and control systems, allowing them to perform complex tasks and maneuvers.

Example: Drones equipped with thermal imaging are used in search and rescue operations. They can quickly survey large areas and identify heat signatures, helping locate missing persons.

Mobile Robots:

Description: These are autonomous or semi-autonomous robots capable of moving in different environments. They use sensors and onboard algorithms to navigate and perform tasks.

Example: Hospitals have started deploying mobile robots for tasks like delivering medications or transporting lab samples. They navigate the corridors, avoiding obstacles and people, ensuring timely and safe deliveries.

Generative AI:

Generative AI creates new content or patterns from existing data. These models can produce unique content - artwork, music, or even realistic human speech. Their capabilities range from generating novel designs to simulating realistic video game environments.

Text Generation:

Description: This pertains to the automatic production of textual content by AI models. These models learn patterns from large datasets and then use this knowledge to produce coherent and contextually relevant text.

Example: OpenAI's GPT series (like the one you're currently interacting with) generates human-like text based on the patterns it learned from vast amounts of textual data.

Image Generation/Manipulation:

Description: Generative models can produce entirely new images or modify existing ones. They can visualize objects, scenes, or entities that don't exist in the real world or make changes to real photos.

Example: NVIDIA's StyleGAN has been used to create incredibly realistic, but entirely fictitious, human faces. There are also apps that can age a person's photo or swap their gender, showcasing image manipulation.

3D Generation:

Description: This involves creating 3D models or environments using AI. These models can be used in simulations, video games, or even real-world prototyping.

Example: Video game developers can use generative AI to automatically produce intricate 3D environments, like forests or cities, without manually placing each element.

Video Generation/Manipulation:

Description: AI can produce entirely new video clips or modify existing footage. This might involve creating realistic scenes, changing elements in videos, or simulating events.

Example: Deepfake technology, a controversial use of AI, can manipulate video footage to make it appear as though a person said or did something they didn't. This technology uses a deep learning model to replace the likeness of one person with another in a video.

Speech Generation:

Description: AI models can generate human-like speech, turning text into spoken words with realistic tone, pitch, and pace.

Example: Google's Duplex can make reservations over the phone by generating speech that's almost indistinguishable from a human caller, even incorporating natural hesitations like "um" and "uh" for realism.

Voice Generation:

Description: Beyond just speech, AI can generate distinct voices, mimicking specific tones, accents, or even specific individuals given enough training data.

Example: Descript's "Overdub" tool allows users to create a unique voice model. Once trained, users can type text, and the system will read it out in the user's own voice, making it useful for content creation without actual voice recording.