What is Artificial Intelligence

Artificial Intelligence (AI) is a branch of computer science that aims to create machines capable of intelligent behavior, mirroring the cognitive functions of the human mind such as learning, problem-solving, and decision-making. The ultimate goal of AI is to improve the efficiency and effectiveness of tasks traditionally requiring human intelligence.

Origins and Early Concepts

- **1940s-1950s**: The groundwork for AI was laid by various scientists, mathematicians, and philosophers who pondered the possibility of machines thinking like humans. Alan Turing's 1950 paper, "Computing Machinery and Intelligence," proposing the Turing Test as a criterion of intelligence, is a foundational moment.

The Birth of AI and Initial Optimism

- **1956**: The term "Artificial Intelligence" was officially coined at the Dartmouth Conference, marking the beginning of AI as a field. Early AI research focused on problem-solving and symbolic methods.

Early Successes and Challenges

- **1960s-1970s**: AI research achieved notable successes in areas like natural language processing, robotics, and games (e.g., chess). However, the field also encountered significant challenges, leading to the first "AI winter," a period of reduced funding and interest due to unmet expectations.

Resurgence and the Second AI Winter

- **1980s**: The development of expert systems, which mimicked the decision-making abilities of human experts, led to a resurgence of interest in AI. Despite commercial success, limitations of expert systems contributed to a second AI winter by the late 1980s.

The Rise of Machine Learning

- **1990s**: The focus of AI research shifted towards machine learning, driven by the realization that machines could learn from data. This period saw the development of algorithms and models that improved over time with exposure to more data.

Breakthroughs in Deep Learning

- **2000s-2010s**: Advances in computational power and data availability facilitated breakthroughs in deep learning, a subset of machine learning. Significant milestones included the success of deep neural networks in image and speech recognition, notably IBM's Watson winning Jeopardy! in 2011, and Google DeepMind's AlphaGo defeating the world champion in Go in 2016.

AI in the Modern Era

- **2020s-Present**: AI is now integral to various sectors, including healthcare, finance, transportation, and entertainment. The development of generative AI models, such as GPT (Generative Pre-trained Transformer) by OpenAI, has opened new frontiers in natural language processing, creative content generation, and beyond. Ethical and societal implications of AI, including bias, privacy, and job displacement, are areas of active research and debate.

AI continues to evolve, with ongoing research aiming to address challenges in generalization, explainability, and ethical use, signaling a future where AI could more profoundly transform human life and industries.

Narrow AI (Artificial Narrow Intelligence, ANI)

Narrow AI refers to artificial intelligence systems designed to perform a specific task or a limited range of tasks. These systems are "narrow" in their capabilities, meaning they are programmed or trained to accomplish particular objectives and do not possess the broader understanding or general intelligence to perform any task beyond those they were designed for. Examples of Narrow AI include speech recognition systems, chatbots, recommendation engines, and image recognition software. Such AI systems are prevalent today, driving much

of the technology we use in our daily lives, from virtual assistants like Siri and Alexa to customer service chatbots and search engines.

General AI (Artificial General Intelligence, AGI)

General AI, on the other hand, refers to a theoretical form of artificial intelligence that can understand, learn, and apply its intelligence to solve any problem with the same level of competence as a human being. AGI would possess the ability to think abstractly, reason, plan, learn from experience, and apply knowledge in different domains. This level of AI would be capable of performing any intellectual task that a human being can. It would have a form of consciousness and self-awareness, making decisions based on reasoning and interpreting the world around it. As of my last update in April 2023, AGI remains a theoretical concept and a goal yet to be achieved in AI research, with significant ethical, philosophical, and technical challenges to overcome.

The distinction between Narrow AI and General AI is fundamental in understanding the current capabilities of AI technologies and their potential future development. While Narrow AI has seen tremendous success and integration into various sectors of society, the creation of AGI represents a future milestone that would mark a profound shift in the field of artificial intelligence.

Key Concepts:

- **Machine Learning (ML):** A subset of AI that enables systems to learn and improve from experience without being explicitly programmed.

    - **Sensing:** Perceiving large amounts of data from sensors in the world and learning how to recognize what's there (images, Sounds, vibrations) Facial recognition, fingerprints, Retinas, self-driving cars recognizing other cars/dogs humans etc. Image Recognition and Analysis.
    - **Predicting:** Programs can use lots of amounts of data of any kind to predict what will happen in the future- Predict financial fraud, disease based on symptoms and clinical lab tests, mechanical failure based on sound a vibration data, even crop yields based on photos of the fields and information on weather etc. Reducing Loss, Personalizing Offerings, improving product performance using better predictions

* Machine Learning is making daily life easier by automating tasks, sorting through data with a level of speed and accuracy superior to that of humans and making connections between data to enhance and personalize the online experience. It is being applied to a wide range of business problems to deliver tangible business value and will continue to transform how people live and work. But its worth mentioning that it still has a way to go before we can trust it. Google Translate is a great example of Machine Learning not quite being fully developed yet in terms of Natural Language Processing.

- **Deep Learning (DL):**

    - A specialized subset of ML that mimics the neural networks of the human brain to process data and create patterns for decision making.

- **Natural Language Processing (NLP):**

- In this fast-moving world, especially if you're a high level leader, the use of NLP's can be transformative in allowing you to add value to your life in a more powerful way if we continue to enable NLP functions to streamline and take the tedious tasks off our plate. This will allow for more time to work on your business, relationships, Health, and Wellness as a few examples which overall, could lead to a more fulfilled life.

- **Robotics:** Robotics is defined as Automated devices that perform physical tasks in the real world.
    - The objective of robots is not to replace humans by mechanising and automating tasks, it is to find ways for humans and machines to be more effective together.

- **Robotic Process Automation (RPA):** RPA is software that automates rules-based actions performed on a computer.
    - RPA is an advanced form of business process automation that can record tasks performed by a human on their computer, then perform those same tasks without human intervention. Essentially, it is a virtual robot copycat.
    - Put simply, the role of RPA is to automate repetitive tasks that were previously handled by humans. The software is programmed to do repetitive tasks across applications and systems. The software is taught a workflow with multiple steps and applications.

Here are a few Use cases for Robotics:

- **Factory Robotics:** Sawyer at Moduform- "Not looking to replace people, looking to create efficiencies and create better more engaging jobs for people to come too.
- **Warehouse Robotics:** Amazon Kiva robots. Allowed amazon to increase inventory by 50% and reduced useless labor such as walking around finding items manually. Their focus on automation is about helping people their jobs not replacing the jobs. They can lift heavier things.
- **Delivery Robotics:** Saviokes Relay robot delivers room service to hotel guests. Can navigate the hotel, get to and from rooms using the elevator, calls the room when it arrives, drops off room service, connects with the guest and then makes its way back. Something interesting that relay does is its able to consider environments to maneuver that's safe for both the guest and the robot.

Rescue Robots: Rescue Robots are being used in disaster situations to send in as First responder support. Disaster robots can help in "making the disaster go away faster". "If I can see what I need to see, I can make good decisions to keep the responders safe." Robots could also be used for security purposes.

AI and Robots:

AI and robots are expected to play a significant role in the future workforce as the collaboration between people and computers increase. Before the COVID-19 pandemic, numerous media reports predicted that robots would replace people on a large scale; however, the arrival of the pandemic reduced the frequency of such reporting. How AI will affect job markets in the future remains an important issue. AI is likely to have a big impact, but not in the way suggested by the pre-pandemic reports that projected massive job losses. Instead, jobs will change and evolve, so that people will work alongside AI and focus their energies on the tasks they do best.

AI and Humans

There are 4 Roles computers can play relative to people.

1) Tool:
- Computers perform the task and Humans monitor the computer Ie. Autocomplete, spreadsheets and cruise control or Tools that connect people to other people (Email, Netflix, Facebook etc.)

2) Assistant:
- Can work without direct attention, take more initiative, more active in solving problems. Ie. Semi-autonomous cars, IBM's Watson technology- Help Drs diagnose medical cases, Chatbots to respond to queries and questions via social media that add suggestive replies to customers, Facebook algorithms for ads.

3) Peer:
- Machines perform similar tasks to people; people must solve some of the cases. LEMONADE (insurance company) AI can pay claims automatically if the claim that was submitted falls into a set of guidelines for that specific claim, but if it doesn't, people will oversee unusual claims.

4) Manager:
- Example: Traffic Lights, Other uninjectable ways we are already using machine to do so are Assigning tasks, coordinating (Who will do which tasks according to data), evaluating, training which are integral roles of a manager

Generative AI: The best way to explain Generative AI is like Dreaming. Dreaming is similar to Generative ai in the sense that our brain stores a mass amount of information, experiences, images, thoughts, people, memories etc. and when we dream, somehow by a magical process that is very difficult to understand, it takes all that "data" or "memory" and stitches bits and pieces into a full experience (dream)

Hallucinations: A Misconception is that a form of generative ai such as Chat GPT is supposed to be saying something that's true, when in fact the models are just trying to predict the next word after all the preceding words. So, in a sense, they aren't trying to say what's true, they are trying to say something that sounds like something one of the human writers upon which the systems were trained would have written (In this case the internet is the data set it's been trained on, so its quite extensive). Generative AI only Hallucinates, most of the time the results happen to be true but sometimes its not. But an extraordinary amount happens to be true because it's been trained on millions of millions of millions of pages of texts on the internet, so a lot of the common ways of saying things are kind of true, but its certainly not specifically designed to always be true.

Thoughts and Reflections on The Future Impact of AI on Society and Business

- AI is transforming industries by optimizing operations, enhancing customer experiences, and creating new business models. It presents opportunities for increased efficiency, but also challenges related to job displacement, privacy, and ethical considerations.
- Successful AI applications center on automating routine work enabling human computer combinations that do revenue increasing work that neither could do alone.
- Then, to have AI Systems that we can trust, we need to have systems that can explain what they are doing (systems that think like we think)
- To reap the benefits of AI, organizations must do more than simply install new AI tools. Rather, as is the case when deploying any new information technology, they must determine the most effective way to combine people and computers.
- What tasks should computers do and what tasks should humans do?
  Machines: Machines are much better at remembering large amounts of information
  People: Interacting flexibly with other people
- We shouldn't just be trying to replace humans wherever we can, we should also be letting the computer systems do things better than any human or computer could do before, which will allow us to do things that were never possible at all before.
- Its important to understand that Machines learn from their own experience with various kind of machine learnings (Chat bots learning from the human response to better respond the way a human would the next time) Allowing Humans to productively work with AI in this manner, the better it will become, and the more opportunities can present themselves.
- AI is intended to Enhance people and our lives, not to replace us. It's time to Elevate Humanity. Don't think of AI as "Artificial Intelligence" think of it as "Augmented Intelligence."
- If we are guided by history, The commercial value of the tools and applications we are developing now, comparably to the ones we've made in the past will be about making new things possible, not necessarily replacing people. So, the myth of AI taking over the world is seemingly getting farther and farther away.
- How do you stay relevant in Life? Those who aren't getting better every day will be replaced with those who are, not just in terms of AI; but everything in life. AI wont replace people, people using AI will replace people who aren't.

What's particularly hard in the current state of AI:

- Adoption times: Ideas to prototypes to production to launch to mass market adoption.
- Interacting with the physical world

- Learning- This Depends on tremendous amount of human common sense. Dissecting where that common sense comes from as a species is needed if we want to progress this Artificially. We Still need to build Systems that Learn like we learn, and this is the current difficulty development in AI is facing.

AI Current Use Cases:

Understanding that AI has already begun to Integrate into our lives over the last few decades is important. Although lots of massive leaps in advancements have started to happen, as society we have already been using Artificial Intelligence in our day to day lives and up until now, you may not have even realized it.
-Siri/Alexa
-Chat GPT
 -Netflix Recommendations
-Smart TV
-Talk to Text
-Google Translate
-Smartphone Autocorrect
- There is also AI for social media content moderation that is used while using social media on a day-to-day basis. These will automatically detect and filter inappropriate and bad behaviour or harmful/offensive content that can result in accounts being banned, or certain features being taken away due to this behaviour.

Unit 1 Exercise and Reflections

- Identify an area in your organization where AI could have a significant impact. Consider both opportunities and challenges.

- Discuss how AI could transform your industry in the next five years. What are the implications for your business?

Unit 2 - Understanding Porter's Business Strategies

Porter's Generic Strategies

Porter's Generic Strategies framework outlines three strategic options for gaining a competitive edge:

Cost Leadership: Being the lowest cost producer in the industry while offering products or services at a standard price.
Differentiation: Offering unique products or services that are valued by customers, allowing the company to charge premium prices.
Focus: Concentrating on a narrow market segment and tailoring products or services to this group better than competitors.

Porter's Strategies in the Digital Age
With the advent of AI, companies can leverage these technologies to reinforce their competitive strategy:

- **Cost Leadership**: AI can optimize operational efficiencies and reduce costs through automation and predictive maintenance.

- **Differentiation**: AI enables the creation of personalized customer experiences and innovative products.

- **Focus**: AI can help identify niche market segments through data analysis and cater specifically to their preferences.

Unit 2 Exercises and Reflections

- Choose one (or all) of Porter's Generic Strategies and brainstorm how AI could support this strategy in your organization.

Unit 3 - Strategic Use of AI in Business

Aligning AI with Business Strategy
The strategic use of AI within the framework of Porter's Generic Strategies enables businesses to enhance their competitive advantage. This section explores how AI initiatives can be aligned with business objectives to drive growth and innovation.

AI and Cost Leadership:
- **Automation of Operations**: Implement AI-driven automation to streamline operations, reduce labor costs, and improve efficiency.

- **Predictive Maintenance**: Use AI to predict equipment failures before they occur, reducing downtime and maintenance costs.

AI and Differentiation:
- **Personalized Customer Experiences**: Deploy AI to analyze customer data and deliver personalized marketing, sales, and service experiences.

- **Innovative Products and Services**: Leverage AI to develop new products and services that meet evolving customer needs.

AI and Focus Strategy:
- **Market Segmentation**: Utilize AI to analyze data and identify niche markets with specific needs.

- **Tailored Solutions**: Develop AI-driven products or services that specifically cater to the identified segments.

Case Studies
- **Retail Giant Using AI for Inventory Management:** A leading retailer employs AI to optimize stock levels, reduce waste, and ensure product availability, enhancing customer satisfaction and driving sales.
- **Financial Services Firm Implementing AI for Fraud Detection:** An innovative financial institution uses AI to analyze transaction patterns in real-time, significantly reducing fraudulent activities and losses.

Unit 3 Activity:

- Select an industry. Identify a business challenge and brainstorm an AI solution that aligns with one of Porter's Generic Strategies. Present Your solution, including its benefits and potential impact on the competitive landscape.

Unit 4: AI Security and Ethics

AI Governance and Management: Unit 4 is for information purposes to start, we will circle back to the workshop after the creation of the roadmap and work on creating Policies for your organization together.

Understanding AI Security Risks

AI systems, like any digital technology, are vulnerable to various security risks that can compromise data integrity, privacy, and operational functionality.

Common Risks:
- **Data Poisoning**: Malicious manipulation of training data to influence AI behavior.

- **Model Theft**: Unauthorized access and theft of AI models.

- **Adversarial Attacks**: Inputs designed to confuse AI models and cause incorrect outputs.

Developing AI Security Policies
Creating robust AI security policies is essential to protect against threats and ensure the integrity of AI systems.

Key Components:
- **Data Protection**: Guidelines for securing training and operational data.

- **Access Control**: Measures to control access to AI systems and data.

- **Incident Response Plan**: Procedures for responding to security breaches.

Ethical Considerations in AI
Ethical AI use is crucial for maintaining trust and alignment with societal values. Key areas of focus include bias mitigation, transparency, and accountability.

Guidelines:
- **Fairness**: Ensure AI systems do not perpetuate or amplify biases.

- **Transparency**: Maintain clarity about how AI decisions are made.

- **Privacy**: Respect user privacy and data rights.

Unit 4: Workshop *To be completed post roadmap development, but it is worth noting for the creation of your roadmap.
- Conduct a workshop to draft AI security policies for your organization. Use breakout sessions to address specific areas like data protection, access controls, and ethical guidelines. Aim to create actionable policies that can be implemented immediately.

Unit 5: Implementing and Managing AI Projects - Best Practices

AI Project Management Best Practices
Successfully managing AI projects requires understanding their unique challenges, such as data dependency and model uncertainty.
Best Practices:
- **Agile Methodology**: Adopt an iterative approach to accommodate changes in AI projects.

- **Cross-functional Teams**: Include diverse skill sets (data scientists, business analysts, IT, etc.) for holistic project perspectives.

- **Continuous Monitoring**: Regularly evaluate AI performance against expected outcomes and adjust as necessary.

Building and Managing AI Teams
The right team composition is critical for AI project success. Focus on creating balanced teams with technical, business, and ethical oversight capabilities.

Considerations:
- **Skills Assessment**: Identify existing skills gaps and plan for training or hiring.

- **Collaborative Culture**: Foster an environment where technical and business teams can collaborate effectively.

- **Ethical Oversight**: Ensure there is expertise in ethical considerations and AI governance.

AI Solution Lifecycle Management
Understanding and managing the lifecycle of AI solutions from development through deployment and maintenance is essential for long-term success.

Lifecycle Stages:
- **Development**: Building and training AI models.

- **Deployment**: Integrating AI solutions into business processes.

- **Maintenance**: Ongoing monitoring, updating, and refinement of AI models.

Unit 5 Exercises and Reflections
- Reflect on a past project (AI or otherwise). How could the adoption of the best practices outlined here have altered the outcome?

- Plan an approach for upskilling your team in AI competencies. Consider both technical and ethical training needs.

Developing an AI Roadmap

Assessing Organizational Readiness
Before embarking on AI initiatives, it's crucial to assess your organization's readiness in terms of infrastructure, data availability, and skills.
Checklist:
- Data Readiness: Do you have access to quality data?

- Infrastructure: Is your IT infrastructure capable of supporting AI technologies?

- Skills: Does your team have the necessary AI skills and understanding, or is there a plan for upskilling and teaching?

Identifying AI Opportunities
Start with identifying business processes or areas that could benefit from AI. Prioritize based on impact, feasibility, and strategic alignment. Work through the Porters Strategies as a reference to help Identify areas of opportunity and concern.
Prioritization Matrix: Create a matrix to evaluate and prioritize AI initiatives based on their potential business impact and implementation feasibility.

Roadmap Development:
Work through the Following Roadmap Worksheet. We will identify the current state of the business, the proposed initiative, and the plan of action and criteria for success.

Roadmap Implementation Planning and Accountability
Develop a phased AI roadmap outlining short-term wins and long-term strategic initiatives. Include timelines, milestones, KPIs, and resource allocation.

Estimated Rollout Phases
- **Phase 1 (0-6 months)**: Quick wins, such as automating repetitive tasks.

- **Phase 2 (6-18 months)**: More complex projects requiring additional data or infrastructure.

- **Phase 3 (18-36 months)**: Strategic initiatives aimed at transforming business models or creating new revenue streams.

Best Practices for Implementing and Managing AI Projects

AI Project Management Best Practices
Successfully managing AI projects requires understanding their unique challenges, such as data dependency and model uncertainty.

Best Practices:
- **Agile Methodology**: Adopt an iterative approach to accommodate changes in AI projects.

- **Cross-functional Teams**: Include diverse skill sets (data scientists, Operations, Front Line team, business analysts, IT, etc.) for holistic project perspectives.

- **Continuous Monitoring**: Regularly evaluate AI performance against expected outcomes and adjust as necessary.

Building and Managing AI Teams
The right team composition is critical for AI project success. Focus on creating balanced teams with technical, business, and ethical oversight capabilities.

Considerations:
- **Skills Assessment**: Identify existing skills gaps and plan for training or hiring.

- **Collaborative Culture**: Foster an environment where technical and business teams can collaborate effectively.

- **Ethical Oversight**: Ensure there is expertise in ethical considerations and AI governance.

AI Solution Lifecycle Management
Understanding and managing the lifecycle of AI solutions from development through deployment and maintenance is essential for long-term success.
Lifecycle Stages:
- **Development**: Building and training AI models.

- **Deployment**: Integrating AI solutions into business processes.

- **Maintenance**: Ongoing monitoring, updating, and refinement of AI models.

- **Review**: Set up a time to review the implemented solutions to ensure continued success.

Exercise and Group Activity for Best Practices:
- Plan an approach for upskilling your team in AI competencies. Consider both technical and ethical training needs.

- Split into teams to outline a plan for an upcoming AI project, including project management approach, team composition, and lifecycle management strategies. Present your plan, highlighting how it addresses the unique challenges of AI projects.

Artificial Intelligence Roadmap Development:

A well-designed roadmap has the potential to catalyze people around the vision for a transformation and the steps necessary to make it real. The roadmap provides a clear and definitive path for moving toward a desired future state. A roadmap should not be considered a static tool, but rather a living document that can be refined and adapted for future needs. It is an effective communication tool for managers and leaders to channel efforts toward achieving the transformation they envision, and then to extend it to new uses in the future. We will be developing this roadmap through the following sections:

The executive summary

The executive summary provides the overview of your roadmap and should concisely encapsulate what you hope to achieve with your AI project and why. In this section of the template, you are required to summarize your plan for using AI in your organization to gain strategic advantage.

Current state
This section requires you to analyze the organization's current use of technology, and its position in relation to competitors. Discuss the organization's strategic positioning in terms of cost leadership, differentiation, or focus.

Proposed initiative
In this section, you are required to detail your proposed initiative for using AI to transform business processes in your organization.

Plan of action and criteria for success
In the last section of your roadmap, describe your plan of action for the implementation of AI in your organization. What steps will you take once you have completed this roadmap?

EXECUTIVE SUMMARY

The executive summary provides the overview of your roadmap and should concisely encapsulate what you hope to achieve with your AI project and why. In this section of the template, you are required to summarize your plan for using AI in your organization to gain strategic advantage. Your summary should include the following:

- Your vision for the AI project

- Organization information, including the relevant products or services it provides.

- Your overall plan for the strategic implementation of an AI technology or a combination of AI technologies

- How the organization can benefit from the use of AI in terms of Porter's three generic strategies: cost leadership, differentiation, or focus

  Note: Your executive summary should be **no more than 500 words**.

 Start Writing Here:

CURRENT STATE

This section requires you to analyze the organization's current use of technology, and its position in relation to competitors. Discuss the organization's strategic positioning in terms of cost leadership, differentiation, or focus. This section covers the following points:

- How innovative technologies are currently deployed in the organization.

- The competitive strategy of the organization (cost leadership, differentiation, or focus)

Note: Your current state analysis should be **no more than 500 words**.

Start Writing Here:

-

PROPOSED INITIATIVE

In this section, you are required to detail your proposed initiative for using AI to transform business processes in the organization. Highlight the following points in your summary:

- The AI technologies or applications that you have chosen to integrate into the organization, and where and how those technologies will be applied.

- Specific tasks and processes in the organization that will benefit from using AI.

- How the use of AI will impact cost leadership, differentiation, or focus and give your organization a competitive advantage (Note that it is generally considered better to focus on only one of the three strategies. However, if you believe some combination of these strategies would be more beneficial in the context of your organization, explain why)

- The technical requirements necessary to roll out the proposed initiative.

- The managerial or leadership requirements necessary to roll out the proposed initiative.

- How your proposed initiative fits in with the organization's IT strategies.

- How your proposed initiative will benefit customers or suppliers

Note: Your proposed initiative section should be **no more than 800 words**.

Start Writing Here:

PLAN OF ACTION AND CRITERIA FOR SUCCESS

In the last section of your roadmap, describe your plan of action for the implementation of AI in your organization. What steps will you take once you have completed this roadmap and this program? Include the following points in your summary:

- An overview of the project plan, with a brief indication of the timelines or milestones

- The skills that you need to build, contract for, or acquire.

- The key stakeholders involved in rolling out your proposed initiative.

- The organizational aspects involved in rolling out your proposed initiative.

- How you intend to manage the initiative

- Any special considerations for rolling out the proposed initiative.

- Any ethical concerns or risks, and how you propose to mitigate them.

- How AI could impact privacy and data protection in your organization. What steps are needed to ensure compliance with regulations?

- The impact that using AI could have on jobs in the organization (Will any jobs be eliminated, or new roles created? Will there be potential to upskill or reskill employees?)

- Criteria by which success will be determined.

- Plans for extending the roadmap to your advantage in the future.

Note: Your Plan of action and criteria for success should be **no more than 800 words**.

Start Writing Here:

IT Partners
Fully Managed IT

"THANK YOU! We hope you and your Business can feel confident in adapting with the greatest technological advancement our world has ever seen. IT Partners Inc. wishes you the best of luck and look forward to assisting you and your needs in the future."