SNIA | Data, Storage & Networking

Introduction to AI and Machine Learning

Live Webinar
August 13, 2025
10:00 am PT / 1:00 pm ET

Today's Presenters

Jayanthi Ramakalanjiyam
Engineering Lead
Celestica
Moderator

Justin Potuznik
Engineering Technologist
Dell Technologies

Tim Lustig
Director, Corporate
Ethernet Marketing
NVIDIA

Erik Smith
Distinguished Engineer
Dell Technologies

The SNIA Community

200
industry leading organizations

2,000
active contributing members

50,000
IT end users & storage pros worldwide

SNIA | DSN A SNIA Community

SNIA | Data, Storage & Networking

What We Do

Drive the awareness and adoption of a broad set of technologies, including:

- Storage Protocols (Block, File, Object)
- Traditional and software-defined storage
- Disaggregated, virtualized and hyperconverged
- AI, including storage and networking considerations
- Edge implementation opportunities and factors
- Storage and networking security
- Acceleration and offloads
- Programming frameworks
- Sustainability

How We Do It

By delivering:

- Expert webinars and podcasts
- White papers
- Articles in trade journals
- Blogs
- Social Media
- Presentations at industry events

Logistics

- The slides are available under the attachments tab at the bottom of your console.
- Questions are welcome!
- Please rate the session and provide feedback!
- Want more sessions like this or other topics, let us know!
  - JOIN US! We meet on Thursday mornings at 11:00 AM eastern.
  - Email dsn-chair@snia.com if you have questions.

SNIA Legal Notice

- The material contained in this presentation is copyrighted by SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

The “AI Stack” Webinar Series

- Building a Strong Foundation for All Experience Levels:
  - Starting from the basics
  - Building steps-by-step
  - Connecting theory to practice
  - Demonstrations
  - Preparing for real-world challenges

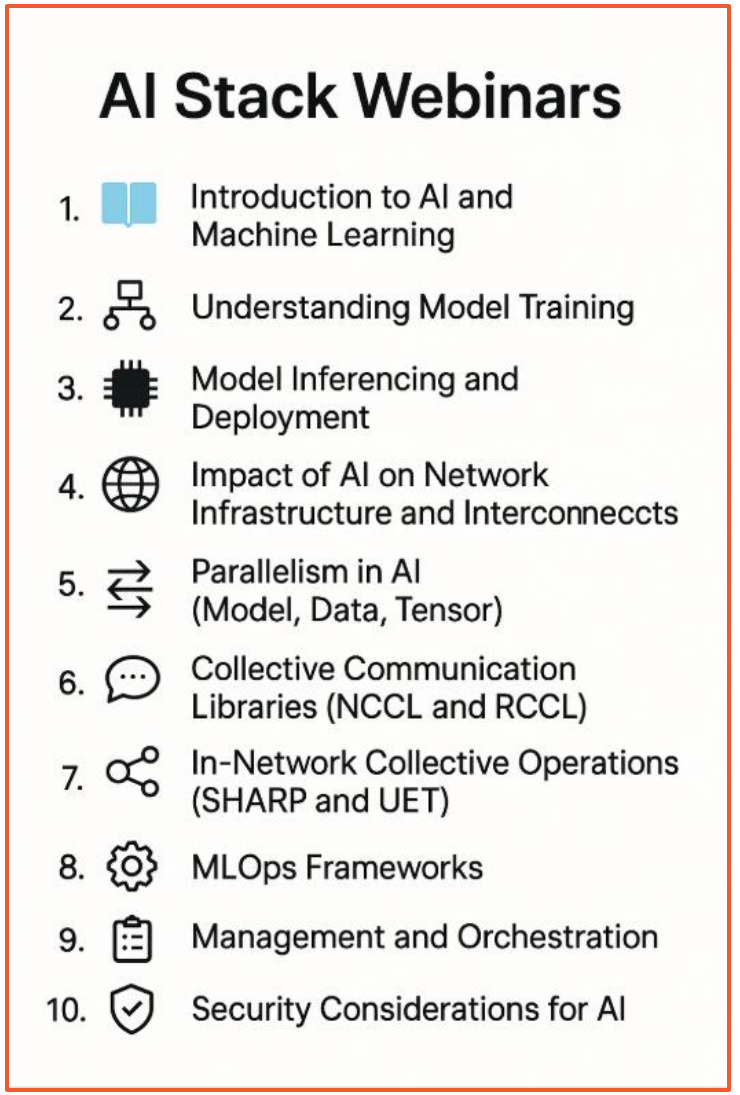

Agenda

01	History and Development Timeline
02	Introduction to AI and ML
03	Learning Techniques
04	Tokens
05	Retrieval-Augmented Generation (RAG)
06	AI Types and Road Map
07	Demo: Training a Tiny Neural Net

History and Development Timeline

Artificial Intelligence History and Development Timeline

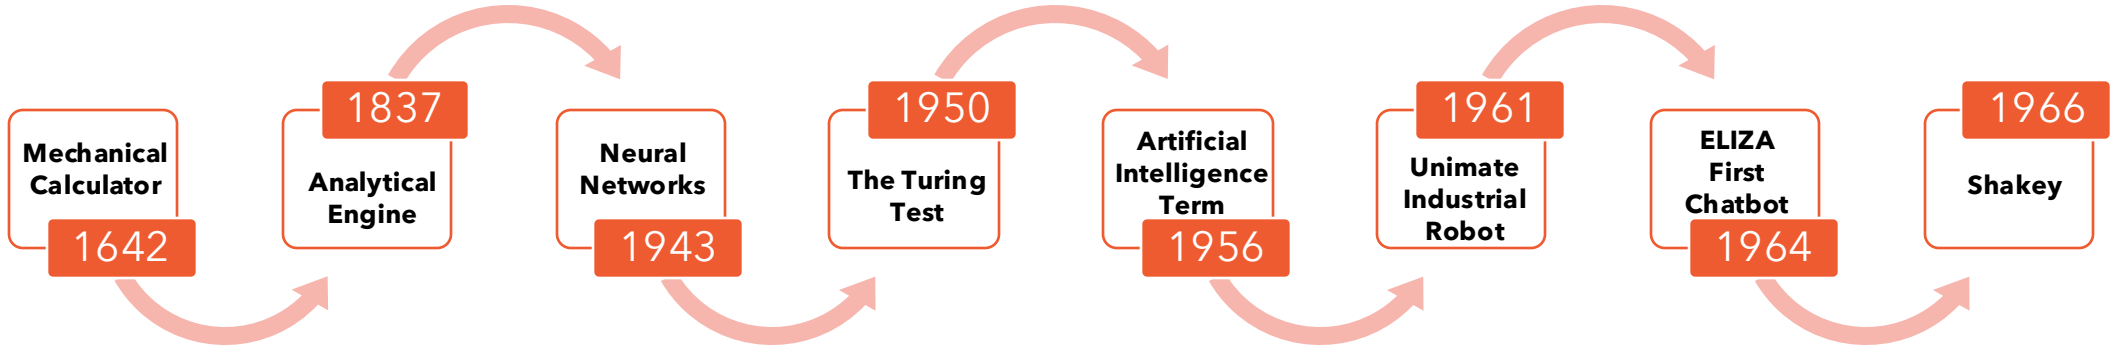

1967 > AI Winter < 1996

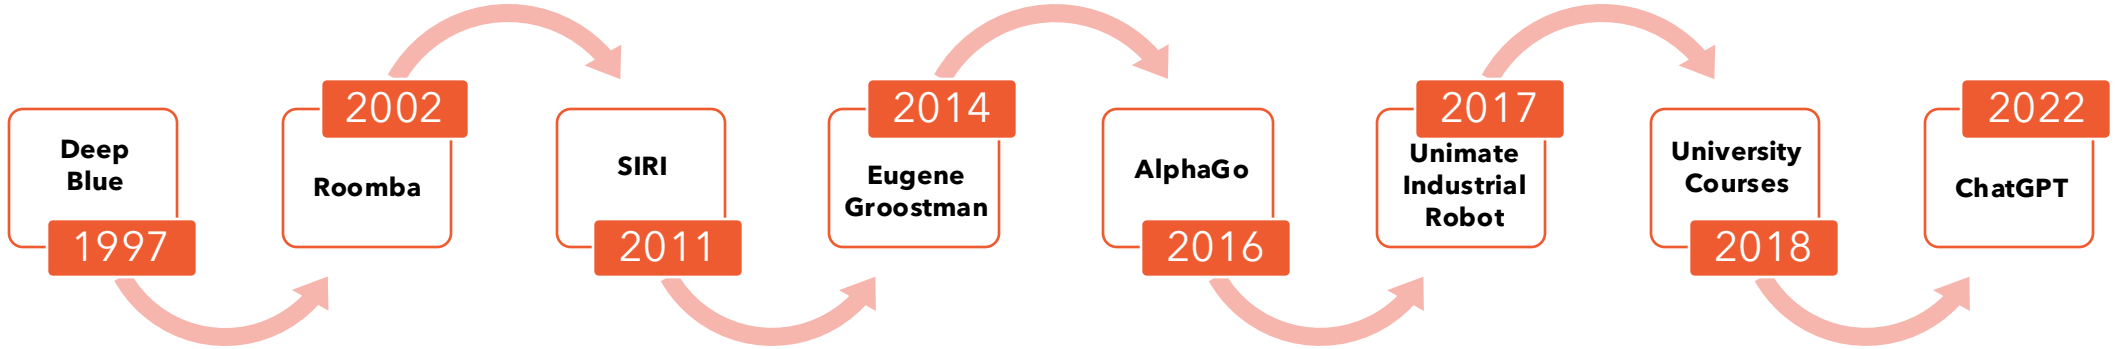

SNIA | DSN A SNIA Community

Introduction to AI and ML

Introduction to AI and ML

- Artificial Intelligence (AI):
  - Broad concept of machines simulating human intelligence.
- Machine Learning (ML):
  - Subset of AI that allows machines to learn from data.
- Deep Learning (DL):
  - Subset of ML using neural networks with many layers.
- Relationship – DL < ML < AI:
  - All part of the same family, but increasingly specialized.
- These technologies power everything from smart assistants to self-driving cars.

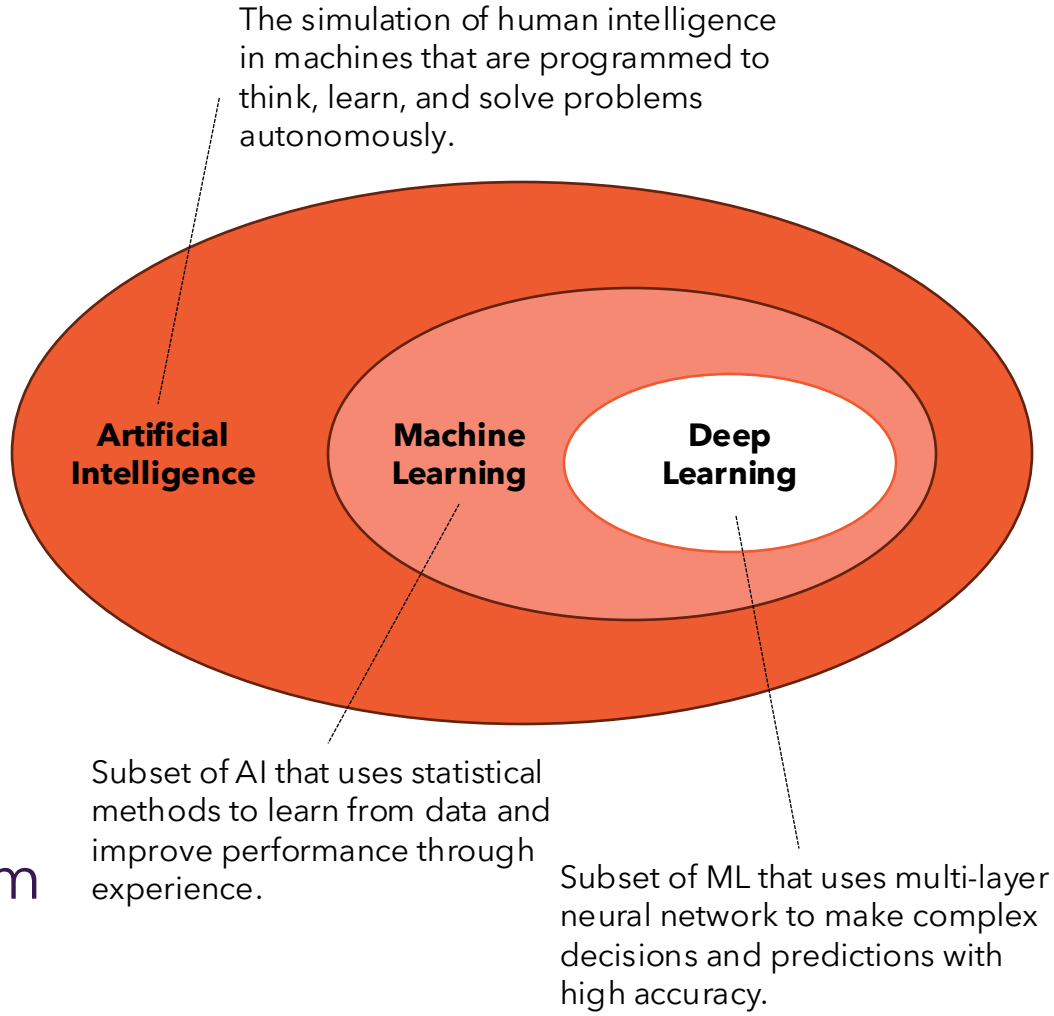

How Does Artificial Intelligence Work?

How does it work?

- Training - Ingesting large amounts of **data (structured and unstructured)** and analyzing the data for correlations and patterns
- Inferencing – looking at patterns to make predictions

Examples of Artificial Intelligence:

- Natural Language Processor (NLP)
  - Interpreting human language for computers to understand. Examples: Siri and Alexa, search engines, spell/grammar checking, translation.
- Computer Vision
  - Enables computers to "see," interpret, and understand the visual world. Examples: Facial recognition, quality control in manufacturing, autonomous vehicles, medical imaging
- Large Language Model (LLM)
  - Primary function is to predict and generate sequences of words. Examples: xAI Grok, Open AI ChatGPT, Google Gemini

SNIA | DSN
A SNIA Community

Learning Techniques

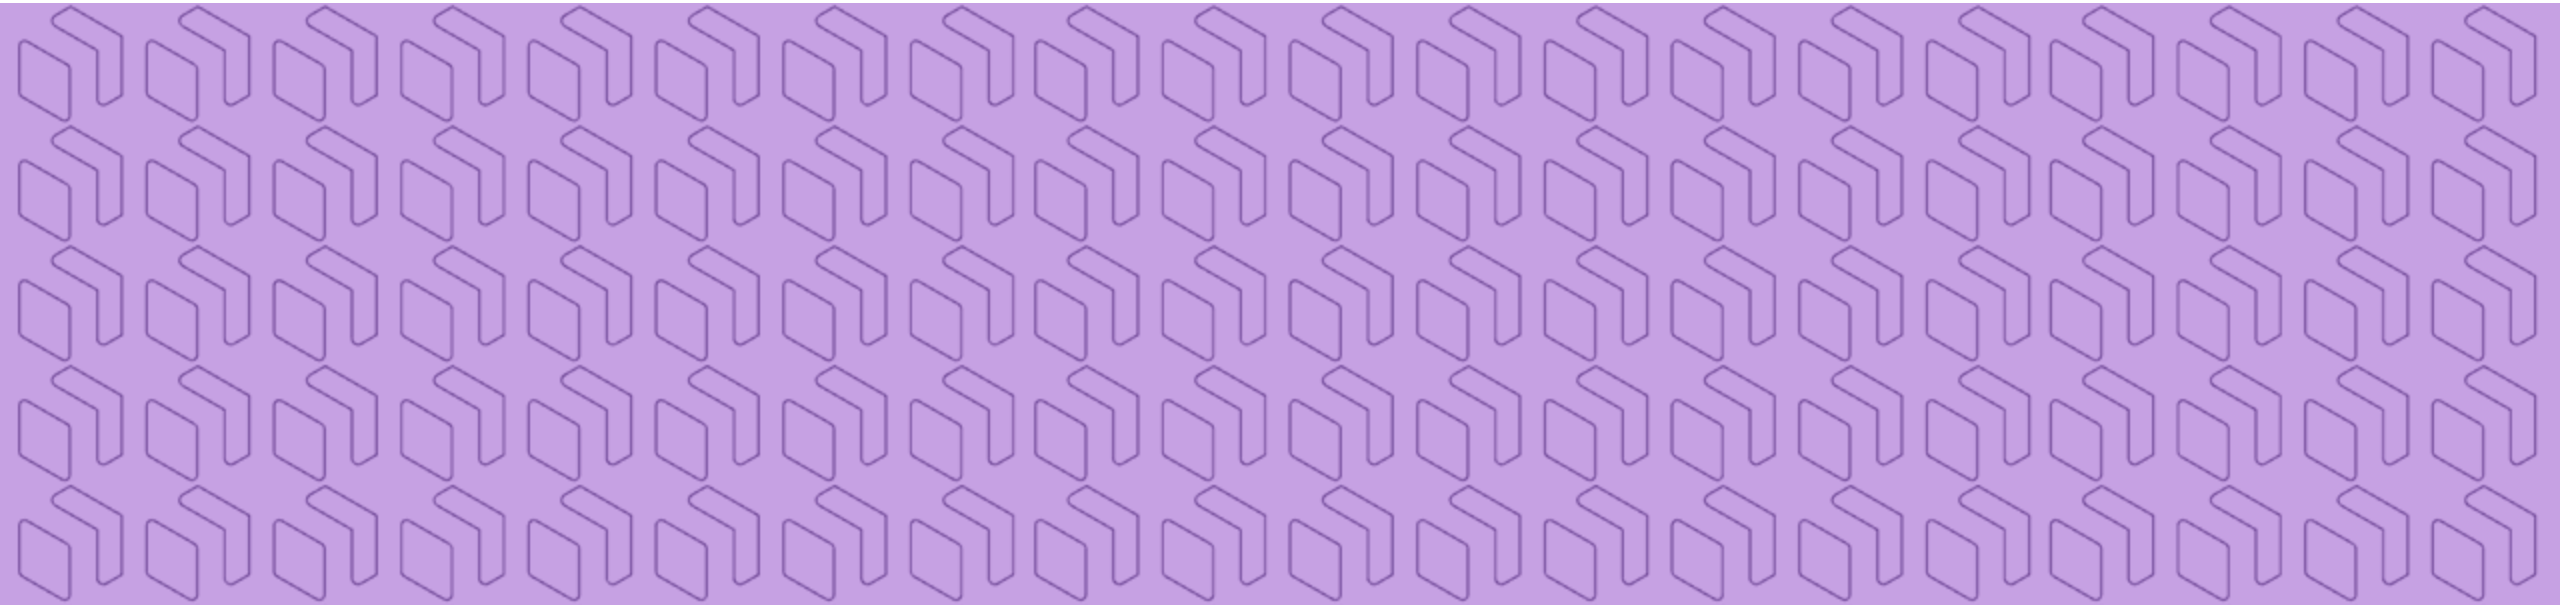

Learning Techniques

- Supervised Learning:

  Algorithms are given labeled inputs for training to learn the relationship between outputs.

- Unsupervised Learning:

  Algorithms learn patterns and structures from unlabeled data.

- Reinforcement Learning:

  An agent learns to make decisions by interacting with an environment and receiving feedback.

Supervised Learning

- Algorithms are given labeled inputs for training to learn the relationship between outputs.

Input: Labeled Data

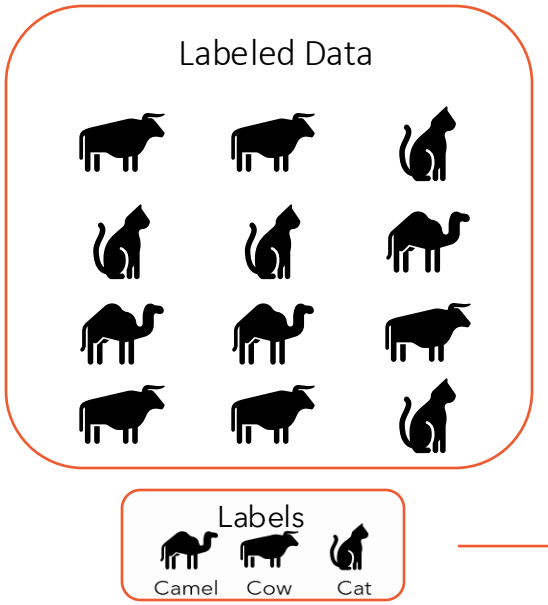

A mixed collection of items with categories given.

Supervised Learning Algorithm

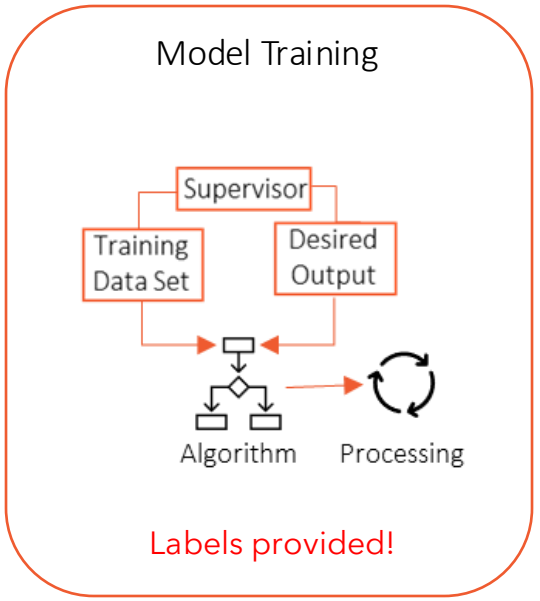

The AI analyzes similarities to find labeled groupings.

Output: Discovered Groups

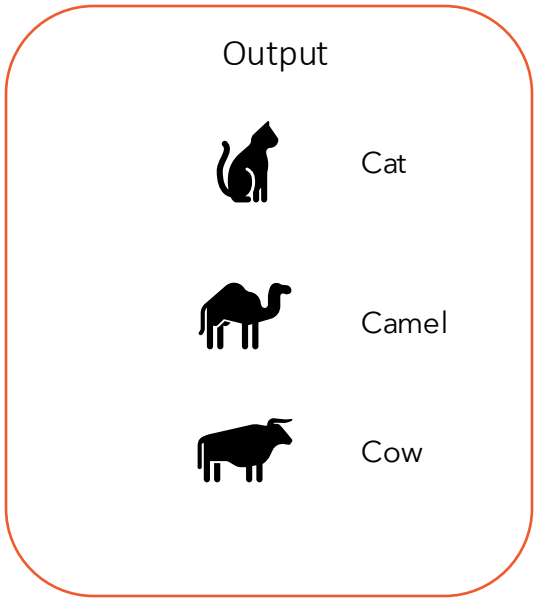

Items are clustered into distinct groups based on labels provided.

Unsupervised Learning

- Algorithms learn patterns and structures from unlabeled data.

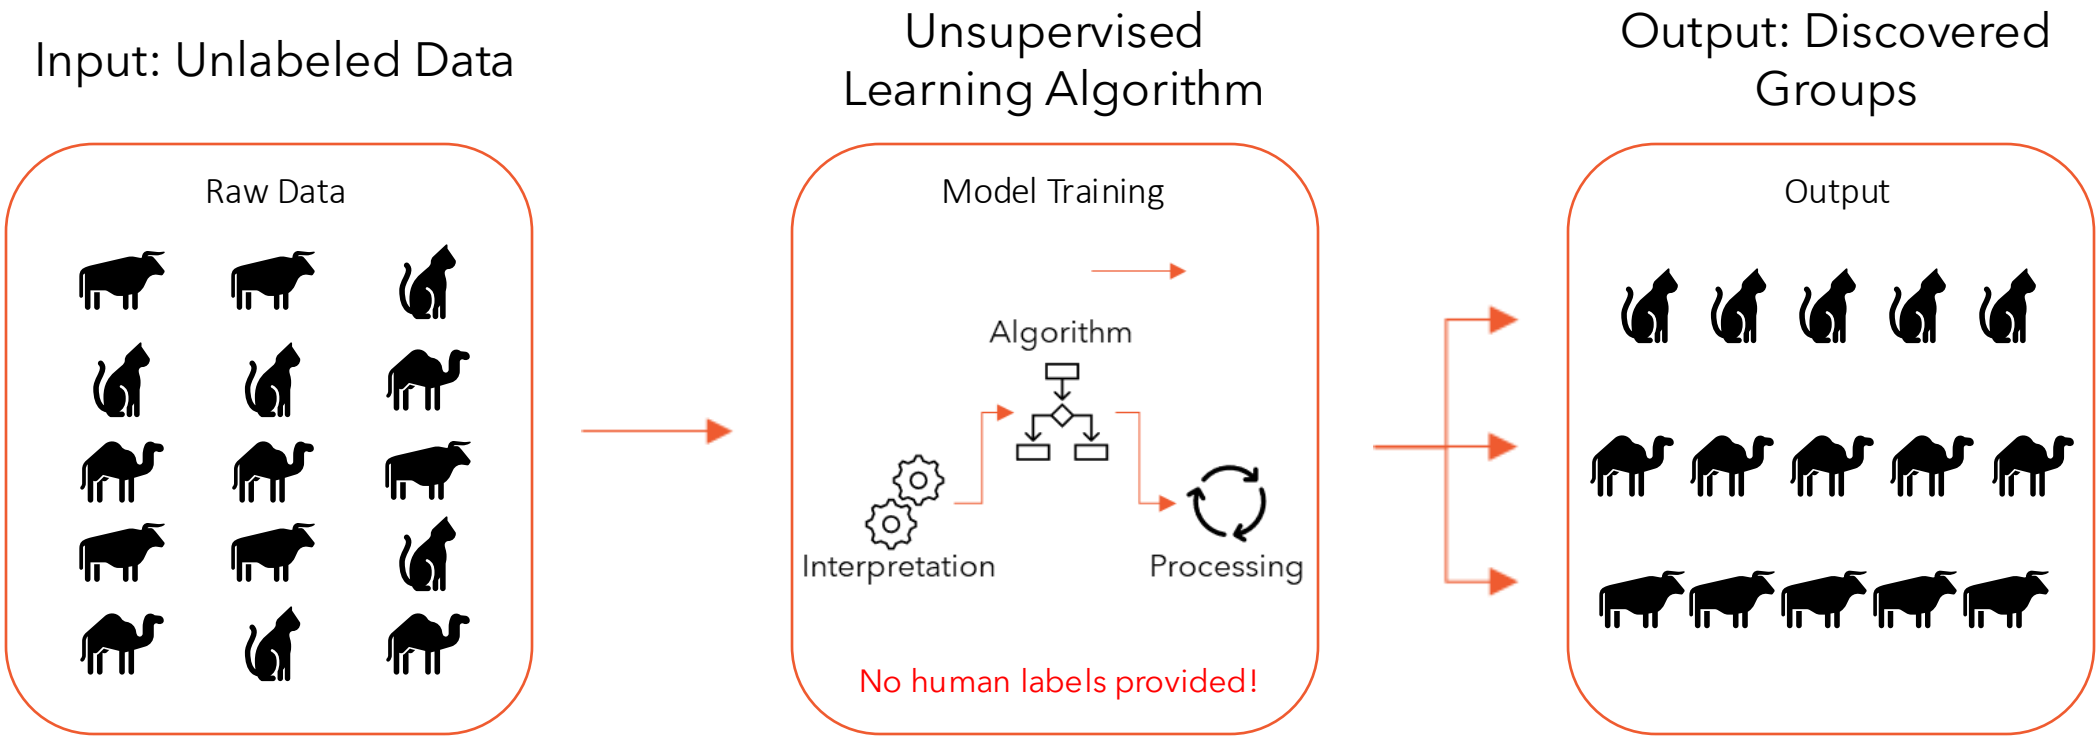

A mixed collection of items with no categories given.

The AI analyzes similarities to find natural groupings.

Items are clustered into distinct groups based on their features.

Reinforcement Learning

- An agent learns to make decisions by interacting with an environment and receiving feedback.

Agent

Decision-making entity

Environment

States

Actions

Rewards

Simulated or physical world

Reinforcement Learning Cycle

Model Training

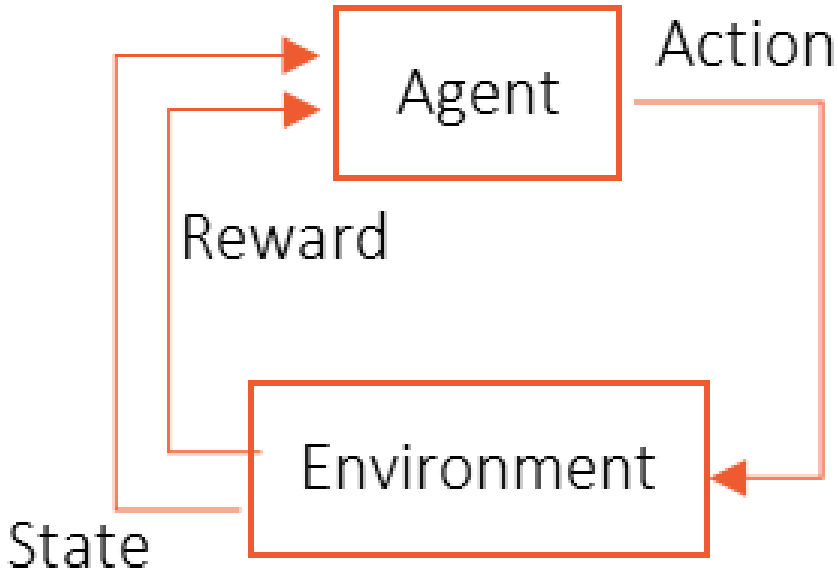

Action
A decision the agent makes in the environment.

Reward
Feedback given by the environment after an action

State
The situation the current agent is in.

Reinforcement Learning

- An agent learns to make decisions by interacting with an environment and receiving feedback.

Agent

Decision-making entity

Environment

States

Actions

Rewards

Simulated or physical world

Reinforcement Learning Cycle

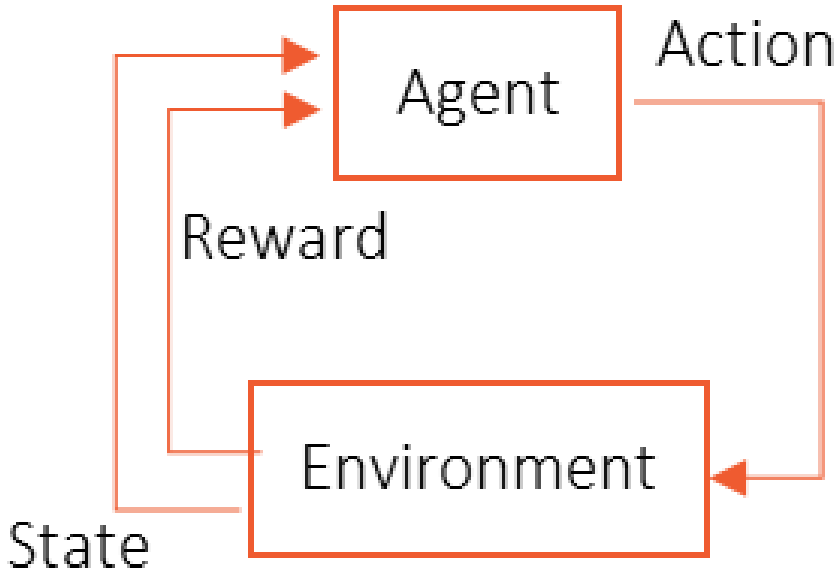

This cycle repeats many times for the agent to learn.

Output: Updated, optimized, Policy

Policy

The learned strategy that achieves the desired results most efficiently.

Tokens

What is a Token?

- A token in AI is the smallest unit of data that a model processes
  - Words are the tokens when processing **Natural Languages**
  - Pixels of the image are tokens when processing **Image Recognition**

- Tokens are the building blocks that AI systems use to understand and generate information.

- Tokenization: breaking down information into manageable pieces

Creating “Tokens” to make Predictions (NLP)

- Natural Language Processing (NLP)
  - Puppies like to eat…

    Creates Tokens

    “Puppies" = 1

    “like" = 123

    “to" = 45

    “eat" = 789
  - [1, 123, 45, 789]
  - “grass", “rocks", “leaves" are very common

Creating “Tokens” to make Predictions (CV)

- Image Recognition - Computer Vision
  1. Turns images into pixels
     1. Color pixels: RGB
     2. Grids
  2. Layers
     1. 1st layer = Outlines
     2. Later layers = Shapes
  3. Patterns
     1. Eyes, ears, nose, tail…
     2. Textures

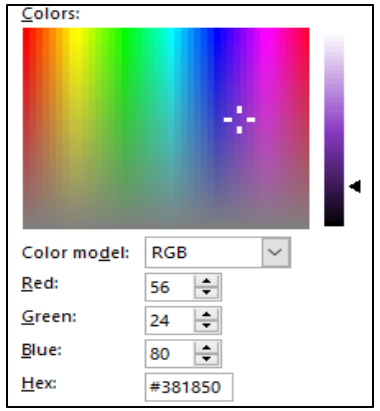

Original Outline

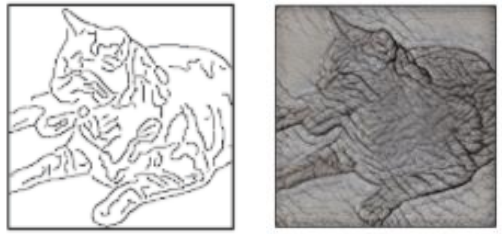

Shapes Texture

Retrieval-Augmented Generation (RAG)

Retrieval-Augmented Generation (RAG)

- Advanced AI architecture that combines language models with a retrieval system.
- Retrieval systems:
  - Online search engine
  - Active database to query
- Why is RAG needed?
  - Fights hallucination – making up facts if the information is not known
  - Provide up-to-date information; knowledge base, products, legal/medical, financial

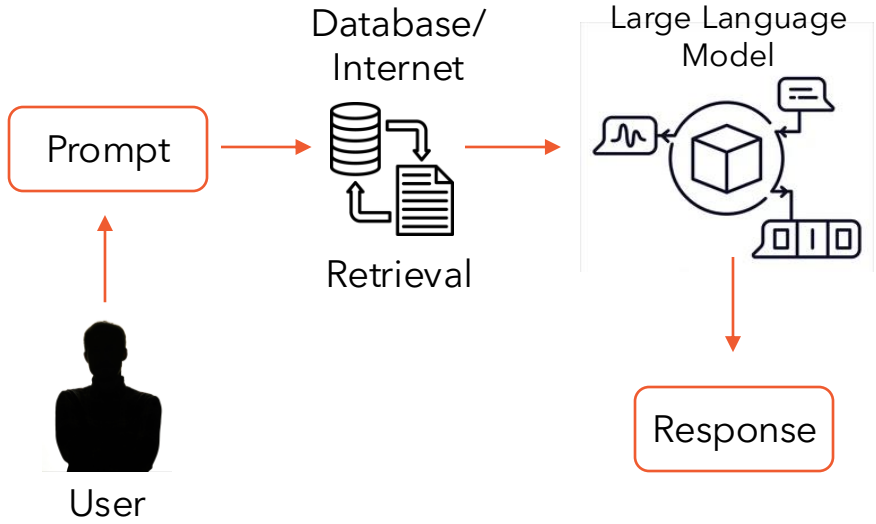

AI Types and Road Map

Types of AI

	Generative	Agentic	Physical
	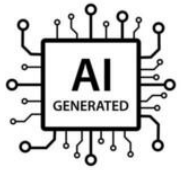		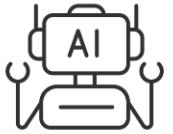
Definition	AI systems that can create new content such as text, images, music, code, or video.	Can use common sense, learn continuously, and reason like a human.	AI systems that perceive, understand, and act directly within the physical world.
Capabilities	Trained on large datasets, using models like transformers (e.g., GPT, DALL·E).	Adaptability and full cognitive mimicry, not just task-specific performance.	Integrates real-time perception, adaptive decision-making, and physical interaction capabilities.
Examples	• ChatGPT (generates text) • DALL·E (generates images) • GitHub Copilot (writes code)	Can learn anything and interact like a human (Data from Star Trek)	Warehouse and factory robots, robotic surgery systems, and autonomous cars and drones.
	Create content	Act autonomously	Interact with objects

Artificial Intelligence – Road Map

	Artificial Narrow Intelligence	Artificial General Intelligence	Artificial Super Intelligence
		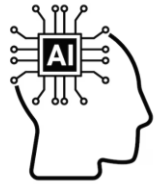	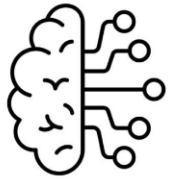
What is it	Designed and trained for one specific task or a limited set of tasks.	Can understand, learn, and apply intelligence across a wide range of tasks .	Surpass human intelligence and have capabilities for problem-solving.
How it works	Learns patterns to perform a single programmed and assigned task.	Can use common sense, learn continuously, and reason like a human.	ASI is theorized to emerge rapidly from AGI through recursive self-improvement.
Key Differentiator	Think of it as a highly skilled tool, not a generalized thinker.	Adaptability and full cognitive mimicry, not just task-specific performance.	Superiority beyond humans, leading to capabilities currently unfathomable to us.
Examples	Voice assistants (Siri), recommandation engines (Netflix), LLMs (ChatGPT).	Can learn anything and interact like a human (Data from Star Trek)	Capable of making scientific discoveries and creating new knowledge.
	Machine Learning	Machine Intelligence	Machine Consciousness

SNIA | DSN A SNIA Community

Demo: Training a Tiny Neural Net: 0 vs 1

- 1957 – Rosenblatt's Perceptron: first trainable neural net
- 1960 – Widrow & Hoff's ADALINE: introduced gradient descent
- 1969 – Minsky & Papert: exposed limits of single-layer nets (e.g., XOR)
- 1974–1986 – Backpropagation emerges and gains traction
- 1989–1990s – LeCun applies backprop to convolutional nets (e.g., LeNet)
- Today – Gradient descent + backprop still power modern neural nets
- Demo Goal – Train a tiny net to answer: Is it a 0 or a 1?

SNIA | DSN
A SNIA Community

Demo

SNIA | DSN
A SNIA Community

Q&A

After this Webinar

- Please rate this webinar and provide us with your feedback
- This webinar and a copy of the slides are available at the SNIA Educational Library snia.org/educational-library
- A Q&A from this webinar, including answers to questions we couldn't get to today, will be posted on our blog at sniablog.org
- Follow us on LinkedIn and X @SNIA

SNIA | DSN
A SNIA Community

Thank You