

DEEP DIVE INTO AI

Understanding, Applying & Advancing Artificial Intelligence



Dr. Md Saef Ullah Miah

Associate Professor, Department of Computer Science

Additional Director, AIUB-IQAC

American International University-Bangladesh (AIUB)



TODAY'S JOURNEY

A structured deep dive from fundamentals to frontier

01



What is AI?

Concepts, history & the AI hierarchy

02



ML Techniques

Algorithms, models & real-world use cases

03



Deep Learning

Neural networks, CNNs, RNNs & Transformers

04



Computer Vision

How machines see and understand images/video

05



Natural Language Processing

Text, speech & language understanding

06



LLMs & Generative AI

Frontier models, prompting & agentic systems

07



Where We Stand (2026)

Global landscape, milestones & what's next

08



Why Learn & How

Roadmap, resources & learning paths

09



AI for CNRD

Naval applications, use cases & integration

01

WHAT IS ARTIFICIAL INTELLIGENCE?

Building machines that think, learn & act intelligently



WHAT IS ARTIFICIAL INTELLIGENCE?

AI is the simulation of human intelligence by computer systems — encompassing learning from experience, reasoning through problems, perceiving the environment, and understanding language.

Artificial Intelligence (AI)

Broad field — any technique enabling machines to mimic human cognition

Machine Learning (ML)

Systems that automatically learn and improve from data

Deep Learning (DL)

Multi-layer neural networks — the engine of modern AI

Foundation Models /

LLMs

Massive pre-trained models powering GenAI

TYPES OF AI

Narrow AI (ANI)

Excels at ONE task. All deployed AI today — chatbots, image classifiers, recommendation engines.

General AI (AGI)

Human-level reasoning across ALL domains. Actively pursued — not yet achieved.

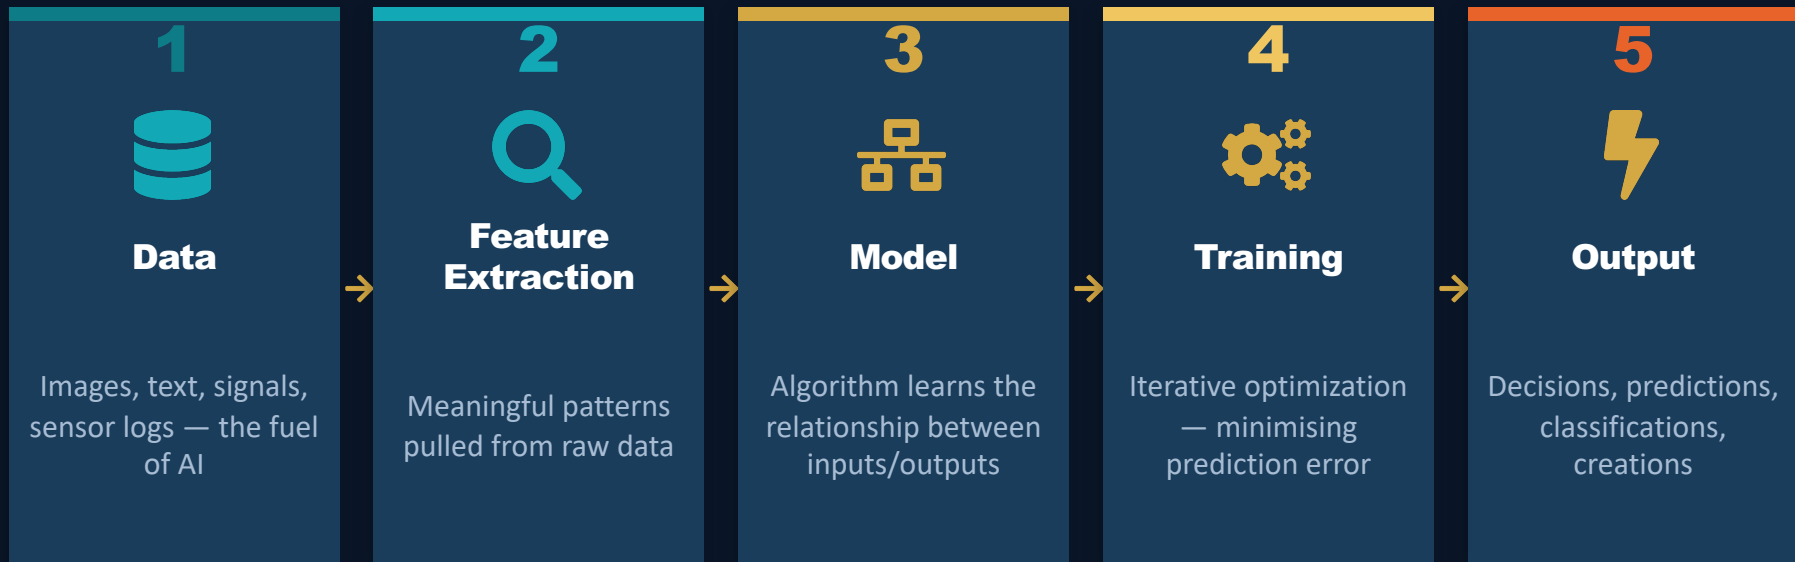
Super AI (ASI)

Far surpasses human intelligence in every respect. Hypothetical — long-horizon goal.

All AI in production today is Narrow AI — yet it is already reshaping industries, militaries & science

HOW DOES AI ACTUALLY WORK?

From raw data to intelligent output — simplified pipeline



Supervised Learning

Learns from labeled data (input→output pairs)

Unsupervised Learning

Finds hidden patterns in unlabeled data

Reinforcement Learning

Learns by reward & penalty — trial and error

These three learning paradigms underpin everything — from spam filters to autonomous naval vessels

02

MACHINE LEARNING TECHNIQUES

Classical algorithms, ensemble models & their real-world use cases



CLASSICAL ML ALGORITHMS — SUPERVISED LEARNING

Learn from labeled data to make predictions



Linear / Logistic Regression

Predict continuous values or class probabilities. Fast, interpretable.
Use case: Ship fuel consumption forecasting, defect probability scoring.



Decision Trees & Random Forests

Tree-based rules for classification & regression. Handles mixed data.
Use case: Equipment failure classification, threat level categorisation.



Support Vector Machines (SVM)

Finds optimal boundary between classes in high-dimensional space.
Use case: Radar signal classification, intrusion detection.



k-Nearest Neighbours (kNN)

Classifies based on similarity to nearest training examples.
Use case: Sensor anomaly detection, pattern matching in sonar data.



Naive Bayes

Probabilistic classifier based on Bayes' theorem — fast & effective.
Use case: Text classification, spam/signal filtering, document tagging.



Gradient Boosting (XGBoost)

Ensemble of weak learners — often best on tabular data. Highly competitive.
Use case: Maintenance prediction, logistics optimisation, risk scoring.

Classical ML remains the workhorse for structured/tabular data — fast to train, easy to explain, production-ready

UNSUPERVISED & REINFORCEMENT LEARNING

Finding structure in unlabeled data — and learning through interaction

UNSUPERVISED LEARNING

K-Means Clustering

Groups data into k clusters by similarity.
Use case: Grouping vessel behaviour patterns, network traffic segmentation.

Principal Component Analysis (PCA)

Reduces data dimensions while preserving variance.
Use case: Compressing sensor data, noise reduction in signals.

Anomaly Detection

Identifies data points that deviate significantly from the norm.
Use case: Detecting unusual ship movements, cyber intrusions, equipment faults.

REINFORCEMENT LEARNING

Agent

The decision-making model

Environment

The world the agent acts in

Reward

Feedback signal — good or bad move

Policy

Strategy learned to maximise reward

Autonomous Navigation

Vessels & UAVs learning optimal routing and obstacle avoidance

Game-Theoretic Warfare

AI learns multi-step tactics through simulated adversarial scenarios

Resource Optimisation

Scheduling, logistics and power management under constraints

RL powers AlphaGo, autonomous drones & robot control — the closest ML paradigm to how humans learn by doing

03

DEEP LEARNING & NEURAL NETWORKS

The architecture revolution that powers modern AI



DEEP LEARNING — NEURAL NETWORKS & CNNs

How layered networks learn rich representations

HOW A NEURAL NETWORK LEARNS

Input Layer

Raw data enters — pixels, sensor values, text tokens

Hidden Layers

Each layer extracts progressively abstract features

Output Layer

Final prediction — class label, value, or probability

Backpropagation

Error signals flow backward to update weights

Gradient Descent

Iteratively minimises the loss function

CONVOLUTIONAL NEURAL NETWORKS (CNNs)

Convolutional Layer

Applies filters to detect edges, textures, shapes in images

Pooling Layer

Downsamples feature maps — reduces computation, retains key info

Fully Connected

Combines all features to produce final class predictions

CNN USE CASES

- ✓ Object detection in UAV imagery
- ✓ Medical image analysis (X-ray, MRI)
- ✓ Ship classification from satellite images
- ✓ Face & biometric recognition

CNNs learn to SEE — they are the backbone of computer vision systems in drones, cameras & surveillance

DEEP LEARNING — RNNs, LSTMs & TRANSFORMERS

Handling sequences, time-series & the architecture behind LLMs

Recurrent Neural Networks (RNNs)

Designed for sequential data — each step depends on previous states. Processes time-ordered data like sensor streams, speech, text.

- ✓ Time-series anomaly detection
- ✓ Speech recognition
- ✓ Predictive maintenance from sensor logs

Long Short-Term Memory (LSTM)

Advanced RNN with memory gates that control what to remember or forget. Solves the 'vanishing gradient' problem — handles long sequences.

- ✓ Sonar signal classification
- ✓ Trajectory prediction for vessels
- ✓ Naval log analysis & forecasting

Transformer Architecture

Uses 'self-attention' — every token attends to every other token simultaneously. Parallelisable, scalable. The backbone of ALL modern LLMs.

- ✓ Powering GPT, Claude, Gemini
- ✓ Document understanding & summarisation
- ✓ Code generation & translation

The Transformer (2017) changed everything — it enabled LLMs and is now the dominant architecture across AI

04

COMPUTER VISION (CV)

Teaching machines to see, interpret and act on visual data



COMPUTER VISION — TECHNIQUES & NAVAL APPLICATIONS

How machines extract meaning from images, video & sensor feeds



Image Classification

Assigns a label to an entire image.
Example: Is this satellite image a warship, cargo vessel or submarine?



Object Detection

Detects & locates multiple objects in one image with bounding boxes.
Tool: YOLO (You Only Look Once) — real-time, fast, widely deployed.



Image Segmentation

Labels every pixel in an image.
Semantic (class) or Instance (object).
Use: Precisely mapping coastlines, separating vessels from sea clutter.



Object Tracking

Follows objects across video frames over time.
Use: Tracking vessel movements, UAV following a target autonomously.



Anomaly Detection (CV)

Identifies unusual visual patterns that deviate from normal.
Use: Detecting structural damage in hull inspections via drone cameras.



Depth Estimation & 3D

Estimates distance and reconstructs 3D scenes from 2D images.
Use: UAV obstacle avoidance, port mapping, underwater robot navigation.

COMPUTER VISION — FRAMEWORKS, TOOLS & MODELS

The practical toolkit for building CV systems today

OpenCV

Open-source CV library — image processing, video analysis, camera feeds. The foundation of most CV pipelines.

YOLOv8 / YOLO11

State-of-the-art real-time object detection. Train on custom datasets in hours. Runs on Jetson edge devices.

PyTorch Vision

Pre-trained models (ResNet, EfficientNet, DETR) + training utilities. Most popular for CV research.

Roboflow

End-to-end platform: annotate data → train model → deploy. Ideal for custom military object detection datasets.

SAM (Segment Anything Model)

Meta AI's foundation model — segments any object in any image with a single prompt. Zero-shot.

NVIDIA Jetson

Edge AI hardware — run CV models at the point of capture (drone, ship camera) without cloud connectivity.

05

NATURAL LANGUAGE PROCESSING (NLP)



Teaching machines to read, understand and generate human language

NLP — CORE TECHNIQUES & TASKS

How machines process and understand human language



Text Classification

Assigns categories to text — sentiment, topic, intent.
Use: Classifying intelligence reports, flagging priority signals in communications.



Named Entity Recognition (NER)

Identifies entities — names, locations, organisations, dates.
Use: Extracting vessel names, ports and dates from intelligence text.



Machine Translation

Translates text between languages automatically.
Use: Translating intercepted foreign naval communications in real time.



Text Summarisation

Condenses long documents into key points automatically.
Use: Summarising lengthy R&D reports, technical manuals, field dispatches.



Question Answering (QA)

Extracts or generates answers from a given document context.
Use: Querying technical manuals, maintenance records without manual search.



Speech Recognition (ASR)

Converts spoken language to text in real time.
Use: Voice commands for ship systems, transcribing radio communications.

NLP bridges AI and human communication — every text-heavy naval workflow has a high-value NLP application

NLP EVOLUTION — FROM RULES TO TRANSFORMERS

How NLP went from pattern matching to human-level language understanding


1950s–90s	Rule-Based NLP Hand-crafted linguistic rules. Dictionaries, grammar parsers. Brittle — breaks on new inputs. Cannot learn from data.
→	
2000–2013	Statistical NLP Machine learning on text features — TF-IDF, n-grams, Naive Bayes, SVMs. Better, but no semantic understanding.
→	
2013–2017	Word Embeddings Word2Vec, GloVe — words mapped to dense vectors capturing meaning. 'King – Man + Woman ≈ Queen'. Semantics emerge!
→	
2017–2019	Seq2Seq + Attention Encoder-decoder with attention mechanism. Revolutionised translation & summarisation. Precursor to Transformers.
→	
2017–Now	Transformers & LLMs 'Attention is All You Need' — self-attention at scale. BERT, GPT, Claude, Gemini. NLP and AI are now one field.

Every NLP technique built on the previous — the Transformer didn't replace NLP, it unified and supercharged it

06

LLMs & GENERATIVE AI FRONTIER SYSTEMS

From language models to reasoning agents



LARGE LANGUAGE MODELS (LLMs) — WHAT & HOW

The technology behind ChatGPT, Claude, Gemini and beyond

An LLM is a Transformer-based neural network trained on trillions of tokens of text — it learns to predict the next token so well that it develops emergent abilities: reasoning, coding, translation, summarisation, and creative writing. Scale changes everything.

HOW LLMs ARE TRAINED

1

Pre-training

Train on massive internet-scale text (billions of pages). Learns language patterns, facts, reasoning structures.

2

Instruction Tuning

Fine-tune on curated instruction-response pairs. Model learns to FOLLOW instructions, not just predict text.

3

RLHF

Reinforcement Learning from Human Feedback — humans rate outputs. Model learns to be helpful, harmless & honest.

EMERGENT CAPABILITIES

- ✓ Reads, summarises & writes documents
- ✓ Reasons through multi-step problems
- ✓ Answers questions from long documents
- ✓ Writes & debugs code in 20+ languages
- ✓ Translates across 100+ languages
- ✓ Generates structured reports & analysis

LLMs did not just improve NLP — they created a new paradigm: a single model that can do almost everything with language

THE LLM LANDSCAPE — KEY MODELS IN 2026

The frontier models shaping AI today

GPT-5.4

OpenAI

Flagship multimodal model — text, vision, voice, real-time reasoning. Powers ChatGPT. Leading coding & reasoning.

Claude Opus 4.6

Anthropic

Safety-focused, long context (500K+ tokens), excellent at analysis, writing & following complex instructions.

Gemini 3.1

Google

1M+ token context window — processes entire books/codebases. Integrated with Google products & search.

Llama 4

Meta (Open)

Open-source — run locally on your own hardware. No API cost. Privacy-preserving for sensitive data.

Mistral / Mixtral

Mistral AI

Highly efficient open models. Mixture-of-Experts architecture — fast inference at lower compute cost.

DeepSeek 3.2

DeepSeek

Strong reasoning model from China. Competitive with frontier models at a fraction of training cost — raises strategic questions.

Open-source models (Llama, Mistral) are strategically important for defence — run on-premise, no data leaves your environment

USING LLMs — PROMPT ENGINEERING & PRACTICAL PATTERNS

How to get the best out of LLMs — the art and science of prompting

PROMPT ENGINEERING TECHNIQUES

Zero-Shot Prompting

Ask directly — no examples. Works for simple, well-defined tasks.

Few-Shot Prompting

Provide 2-5 examples in the prompt. LLM learns the pattern & replicates.

Chain-of-Thought (CoT)

Ask the model to 'think step by step'. Dramatically improves reasoning on complex tasks.

Role Prompting

Assign a persona: 'You are a naval engineer...' — shapes tone, depth and domain focus.

RAG (Retrieval-Augmented Generation)

Connect LLM to your own documents. Ask questions about internal manuals, reports without retraining.

LLM USE CASES FOR CNRD

R&D Literature Review

Feed 50 research papers — get a synthesised summary, key findings and gaps in minutes.

Technical Documentation

Generate maintenance manuals, system specs and SOPs from engineer notes automatically.

Code Generation & Review

Write Python/MATLAB/C++ code for signal processing, simulation and data pipelines.

Intelligence Summarisation

Summarise field reports, signals logs and intelligence briefs into structured outputs.

Training Simulations

Generate realistic scenario scripts, quiz questions and training materials on demand.

Start with RAG on internal documents — it is the fastest path from zero to a working AI tool with real CNRD value

AGENTIC AI — THE NEXT FRONTIER

LLMs that plan, act and complete complex tasks autonomously

An AI agent is an LLM equipped with tools — web search, code execution, file access, API calls — that can autonomously decompose a goal into steps, execute them sequentially, observe results and adapt. It works like an AI employee, not just a chatbot.



REAL-WORLD AGENTIC AI SYSTEMS

Claude Code

Autonomous coding agent — reads your codebase, writes, runs, tests and fixes code end-to-end.

AutoGPT / CrewAI

Multi-agent frameworks — multiple AI agents collaborate, each with specialised roles, to complete complex projects.

OpenAI Operator

Browses web, fills forms, books tickets — completes real-world tasks in browsers autonomously.

Military AI Agents

DARPA, Pentagon deploying agents for intelligence analysis, logistics planning, and decision support at speed.

Agentic AI is not science fiction — it is already automating knowledge work. CNRD can deploy agents for R&D tasks today

WHERE ARE WE TODAY IN AI?

The current state of AI development — December 2025

\$500B+

Global AI investment

10T+

Parameters in leading
frontier models

2B+

Active users of AI tools
worldwide monthly

75%

Fortune 500 companies
with AI in core ops

2017

Transformer — 'Attention is All You Need'. Changes everything in NLP and beyond.

2020

GPT-3 (175B params) shocks the world with few-shot learning. AI enters mainstream consciousness.

2022

ChatGPT — 100M users in 60 days. Fastest product adoption in history. GenAI era begins.

2023

GPT-4, Claude, Gemini, Llama — multimodal, multi-capable AI. Open-source models emerge.

2024

AI agents, o1 reasoning models, real-time voice/video AI. 50+ nations adopt sovereign AI strategies.

2025

AI passes bar exams, PhD benchmarks, medical boards. AGI debate intensifies. AI in defence operations globally.

2026

Agentic AI in enterprise workflows. AI co-pilots in military ops. Multimodal autonomous systems.

We are no longer asking 'Can AI do this?' — we are asking 'How fast can we deploy it?'

WHY SHOULD WE LEARN AI?

The strategic imperative — especially for a Naval R&D institution



Military Necessity

Modern navies integrate AI in ISR (Intelligence, Surveillance, and Reconnaissance), autonomous vessels, torpedo guidance & decision support. Falling behind means losing strategic edge.



Accelerate R&D

AI compresses research timelines. Months of simulation done in hours. AI-assisted design and testing transforms indigenous capability.



UAV & Autonomy

UAVs can evolve from remotely piloted to AI-guided — enabling beyond-LOS surveillance, threat detection & swarm operations.



Data Intelligence

Naval operations generate enormous sensor data. AI transforms this data flood into actionable intelligence in real time.



Workforce of the Future

Engineers who understand AI are the most valuable talent globally. AI literacy is now a prerequisite for competitive naval R&D.



Indigenous Capability

Relying on imported AI black-boxes creates strategic dependency. Building indigenous AI gives Bangladesh Navy full control & security.

WHAT TO LEARN — SKILL ROADMAP

A progressive path from foundation to frontier



LEVEL 1 FOUNDATION

- ✓ Python basics (NumPy, Pandas, Matplotlib)
- ✓ Math: linear algebra, probability, statistics
- ✓ Core ML: regression, classification (scikit-learn)
- ✓ Jupyter Notebooks / Google Colab
- ✓ Complete fast.ai Part 1 or DeepLearning.AI

Est. time: 3–6 months



LEVEL 2 INTERMEDIATE

- ✓ Deep Learning: CNNs, RNNs, Transformers
- ✓ PyTorch / TensorFlow
- ✓ Computer Vision: YOLO object detection
- ✓ NLP: text classification, summarisation
- ✓ Prompt engineering & LLM API usage

Est. time: 6–12 months



LEVEL 3 ADVANCED

- ✓ Fine-tuning LLMs & RLHF
- ✓ Building AI agents (LangChain, CrewAI)
- ✓ Edge AI: Jetson Nano, FPGA deployment
- ✓ MLOps: production pipelines, monitoring
- ✓ Domain AI: naval signals, sonar, autonomous systems

Est. time: 12–24 months

HOW TO LEARN AI — RESOURCES & TOOLS

The best platforms, courses and tools available today

BEST LEARNING PLATFORMS

fast.ai

Top-down practical approach — build first, theory later. Free.

DeepLearning.AI

Andrew Ng's structured courses — industry gold standard.

Kaggle

Real datasets, competitions, free GPU notebooks. Essential.

Hugging Face

Models, datasets, and state-of-the-art NLP/CV tutorials.

Andrej Karpathy

YouTube: build GPT from scratch — the best DL education online.

Coursera / edX

Certificates from MIT, Stanford, Google, DeepMind.

TOOLS TO START WITH

Python

Universal AI language — free, vast ecosystem

Google Colab

Free GPU/TPU in browser — no setup needed

PyTorch

Most popular research & production DL framework

Ollama

Run LLMs locally — privacy-preserving, offline

LangChain

Build AI agents and RAG pipelines quickly

Roboflow / CVAT

Annotate images, train CV models, deploy fast

Best approach: Pick ONE problem you already care about — build a model around it. That first project teaches more than 100 lectures.

AI INTEGRATION — NAVAL USE CASES



Intelligent UAV Systems

Upgrade UAVs with YOLOv8 object detection, autonomous target tracking, obstacle avoidance & AI mission planning for coastal surveillance.



Maritime ISR & Threat Detection

CV models on radar/sonar feeds detect surface vessels, submarines & anomalies in real time — no human-in-loop required for classification.



Predictive Maintenance

ML models monitoring ship system sensors predict component failures before they occur — reducing downtime and avoiding catastrophic breakdowns.



Cybersecurity & Anomaly Detection

AI-powered intrusion detection for naval networks. Identify abnormal traffic, unauthorised access and electronic warfare signatures instantly.



AI-Assisted R&D (LLM + Agents)

LLMs parse technical literature, suggest design improvements & generate documentation. AI agents run simulations and analyse results autonomously.



SIGINT & Signal Classification

ML models for automatic signal classification, spectrum analysis, and EW detection across maritime frequency bands using NLP + CV techniques.

KEY TAKEAWAYS & CALL TO ACTION

1

AI is infrastructure, not a trend

ML, CV, NLP & LLMs are now the building blocks of every advanced defence system — like electricity, you cannot afford to opt out.

2

The learning path is clear

Foundation → ML techniques → Deep Learning → CV & NLP → LLMs → Agents. Each step builds on the last. Start today.

3

CNRD is uniquely positioned

Existing UAV programs, R&D culture and technical talent make CNRD ready to absorb and deploy AI across all mission areas.

4

Build small, ship fast

Pick ONE problem — train a CV model on UAV footage, build a RAG system on technical manuals. Real projects beat 1,000 lectures.

5

Indigenous AI is a national priority

Bangladesh Navy's AI systems must be locally built and controlled — not foreign black boxes. Sovereignty depends on it.

The future belongs to those who build it. CNRD has the talent. AI gives you the tools. The time is now.

THANK YOU

saef@aiub.edu

Questions & Discussion

"The best way to predict the future is to invent it." — Alan Kay