ARTIFICIAL INTELLIGENCE AND NEURAL NETWORKS

Lecture 1a – Introduction to Artificial Intelligence

Chizhi Chris ZHANG

zhangchizhi@ciomp.ac.cn

Advanced Computing and Digital Technology Research Center

University of Chinese Academy of Sciences

Spring 2026

Today’s Question

Main question

What is AI in a way that is actually useful for this course, and why did learning from data become the center of modern AI?

A scene students already know

Phone unlock, route planning, recommendation, translation, spam filtering, and search ranking all feel ordinary now. They do different jobs, but they share one theme: the system is making judgments from patterns.

What to avoid from the first lecture

We will not treat AI as magic, and we will not reduce it to buzzwords. The useful questions are: what task is being solved, what data teaches the model, and what kind of mistake matters?

AI Is Larger Than One Tool

What we mean by AI here

Artificial intelligence is the broad effort to make machines perform tasks that usually require some form of perception, learning, reasoning, decision, or action.

What that can include

Understanding handwriting, spotting unusual behavior in data, recommending a film, helping read a medical image, or planning several steps toward a goal.

What AI is not. It is not equal to robots, not equal to ChatGPT, and not equal to any one software package.

The Family Map

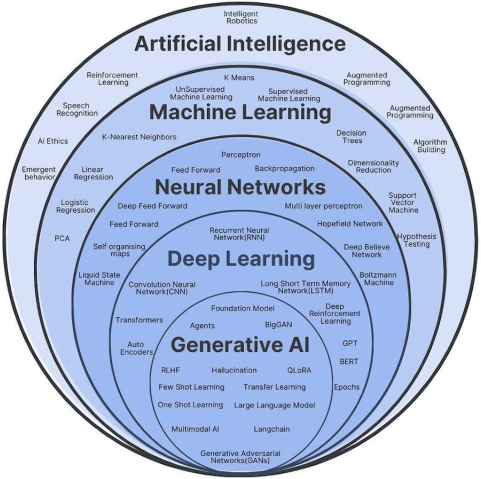

AI is the broadest layer. Machine learning is one route inside AI. Neural networks are one major family inside machine learning. Deep learning is the multi-layer version of that family.

Read it from outside to inside. The closer you move toward the center, the more specific the method becomes.

A concrete reading

ChatGPT belongs to several layers at once: AI, machine learning, neural networks, deep learning, and generative AI.

Automation Is Not Always AI

Where automation ends and learning begins

Ordinary automation follows explicit rules that humans wrote in advance. AI becomes more relevant when the system must learn from examples, adapt to variation, or generalize to cases nobody wrote down one by one.

Simple comparison

An `if-else` script that checks whether a score is above 60 is automation. A model that learns from past student behavior which students may need support next week is closer to AI.

Why students confuse these. Both can look intelligent from the outside. The difference is where the judgment comes from: fixed rules or learned patterns.

Why People Suddenly Notice AI

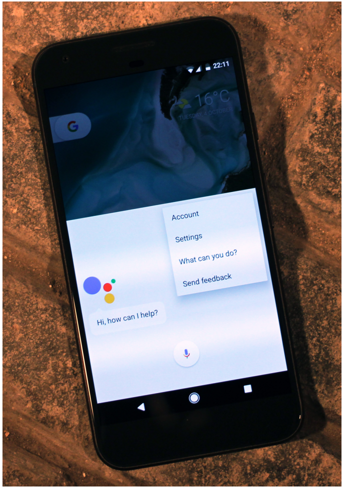

AI became visible to the public when it stopped living only in labs and started sitting inside ordinary screens.

What people actually noticed was not a new formula. It was the feeling that a normal device could suddenly listen, rank, summarize, and answer in one flow.

A familiar moment

You speak to a phone, it hears the words, interprets the request, ranks possible actions, and speaks back in seconds. The novelty is not just the model. It is the closeness of the interface.

Promise and Caution

Why people are excited. AI can notice patterns faster, reduce routine work, and support choices in situations where the information load is too large for a human to scan comfortably.

Low-stakes examples

Captions, translation, recommendation, and search all feel useful because they reduce friction in daily life.

Why people are cautious

The same smooth interfaces can hide weak evidence, biased data, or misplaced trust. A confident answer is not the same thing as a justified answer.

A course habit. Whenever a system looks impressive, ask what kind of error would be costly and who would absorb that cost.

A Day Full of Small AI Judgments

A familiar chain

In one ordinary day, a student may unlock a phone with face recognition, get a route recommendation, receive a spam decision in email, watch a ranked video feed, and use translation on a sign or document.

Why this page matters

These systems do not feel like one dramatic robot moment. They feel like many tiny judgments quietly folded into everyday life.

The useful question

If AI is already embedded in small routine choices, then understanding AI is not only for specialists. It becomes part of ordinary digital literacy.

From Rules to Learning

Input + Program = Output

Traditional programming. Rules + data produce outputs. Humans explain the logic in advance, and the machine executes it.

Input + Output = Program

Machine learning. Data + answers produce a model. Instead of listing every rule, we let the algorithm discover patterns from examples.

Why this shift mattered. Many real tasks are too large, too noisy, or too unstable for a human-written rule-book.

Why Rules Break

The world changes. Speech changes across accents. Images change with lighting and angle. Human behavior changes with time, money, habit, and context. The same fixed rule often ages badly.

A simple story

A fraud pattern that worked last month may fail next month because people changed tactics. A recommendation rule that worked last semester may feel outdated after a platform redesign.

Why learning helps. Learning does not remove uncertainty. It gives us a way to update judgments from evidence instead of rewriting the whole rule set by hand.

A Late-Payment Story

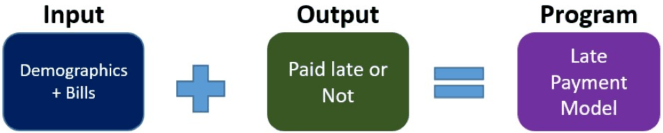

What makes the task realistic

Task. Suppose a company wants to predict which invoices are likely to be paid late.

Why rules struggle. Several clues matter together: amount, season, customer history, and timing. No single clue is enough.

Why this works well in class. It is easy to picture, but messy enough to show why learned patterns beat a short rule list.

What Learning Means

Minimal formula

$$\hat{y} = f_\theta(x)$$

Read it in plain language. Give the model an input x, let it make a prediction $\hat{y}$, and adjust the parameters θ so later predictions become less wrong.

The simple story behind the symbols. Machine learning is repeated correction. The model predicts, compares with reality, and updates itself.

Why one small formula is enough here. At this stage, students do not need a forest of equations. They need the picture: input, prediction, feedback, update.

Data Is the Hidden Teacher

Why data matters so much. The model only knows what the data lets it see. If the examples are narrow, biased, noisy, or mislabeled, the learned pattern can be narrow, biased, noisy, or wrong.

A classroom analogy

If you only teach from one kind of example, students may do well on homework that looks similar and struggle the moment the question changes shape.

The first serious warning

Many apparent “model problems” are actually data problems wearing a model costume.

Why this matters later. The next lectures on data and neural networks only make sense if this point is already clear.

Evaluation Begins With Task Definition

One domain, different tasks. “Will this device fail tomorrow?” and “How many hours remain before failure?” sound related, but one is classification and the other is regression.

Why students lose marks here. If the task is vague, then the model choice, metric, and interpretation all become vague as well.

A graduate-level habit. Before asking whether the model is good, ask whether the question was framed correctly.

This matters for coursework too. Good marks usually come from a well-defined problem, not from the most fashionable model name.

Four Common Task Shapes

A simple first map. Many introductions to machine learning become easier once we sort tasks by the kind of answer they are trying to produce.

Classification	**Regression**	**Clustering**	**Reinforcement**
Choose a label	Estimate a number	Find hidden groups	Choose actions over time

Why this map helps. Students stop trying to memorize method names first and start by asking: what kind of answer is the system meant to produce?

Classification Means Choosing a Label

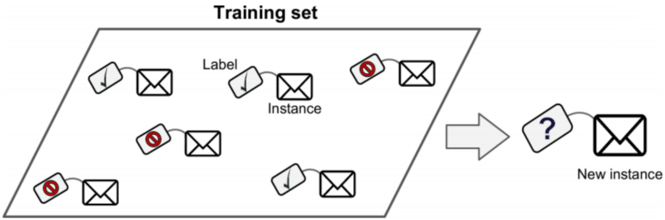

Typical question. Is this email spam or not? Is this image normal or abnormal? Is this applicant likely to repay or not?

Classification is about categories, not quantities.

Why students usually get it quickly. Once the answer is a label, the whole task becomes easier to picture.

Regression Means Estimating a Number

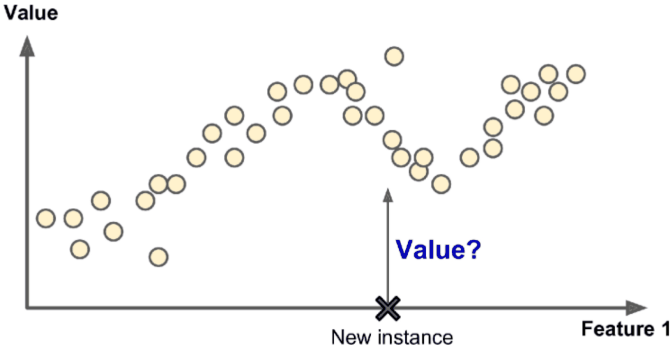

Typical question. What will the temperature be tomorrow? How long will a battery last? What price is reasonable for this house?

Regression is about amount, not category.

The easy way to remember it. If the answer should be a number, students should first think about regression.

Clustering Means Finding Structure

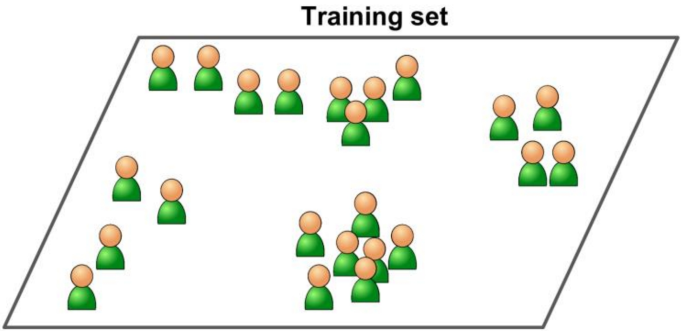

Typical question. If nobody gave us labels, do these samples naturally fall into groups?

Clustering is useful in early exploration, customer segmentation, and research settings where labels are scarce.

This comes up a lot in research. In many real settings, the expensive thing is not data collection. It is trusted labeling.

Reinforcement Learning Means Acting

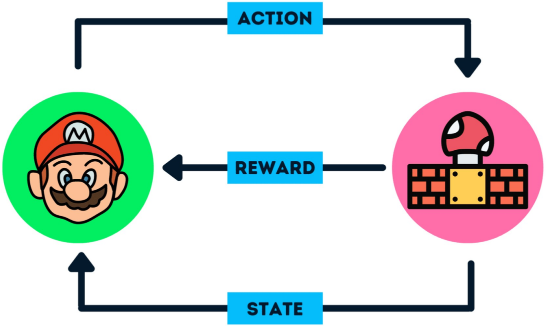

Typical question. What action should the system take now if this step changes what happens next?

Games, robot control, navigation, and other sequential choices fit this shape.

What makes it different

The target is not just a right label. It is a good long-term strategy.

One App Usually Mixes Several Jobs

A phone example. One shopping or music app may recommend items, rank search results, detect unusual behavior, and estimate when you are most likely to return.

What students should notice. Real products rarely use one neat AI task in isolation. They combine several small tasks behind one smooth interface.

Why this matters. It prevents the mistaken idea that AI systems arrive as one giant all-purpose model.

Recommendation Feels Personal

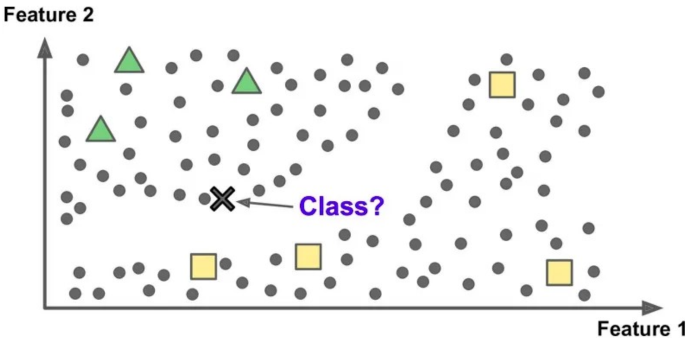

Figure 1-11. Semisupervised learning

The system tries to place similar users or similar items near one another in a learned pattern space.

What feels personal is often just a well-trained similarity judgment running very quickly.

A familiar scene

If two songs attract listeners with similar behavior, the platform treats that overlap as evidence even when the songs do not look similar on the surface.

Rare Events Are Often the Point

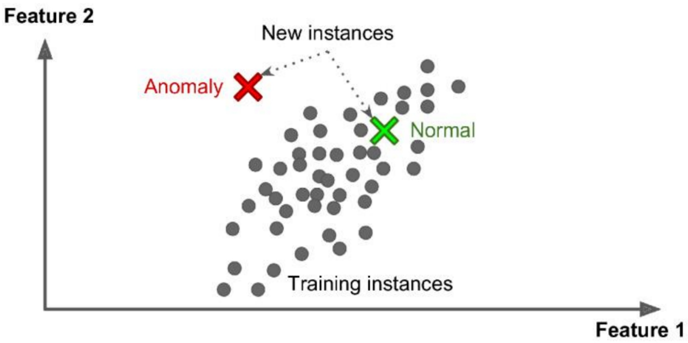

Figure 1-10. Anomaly detection

Sometimes the system is not searching for the most common pattern. It is searching for the unusual case that deserves attention.

Examples

Fraud, unusual machine behavior, strange measurements, or suspicious account activity.

Why it matters. The important event may be rare, but the cost of missing it can be high.

Forecasting Helps Preparation

Forecasting is not fortune telling.
The goal is usually better preparation, not perfect certainty.

Ordinary examples

Demand planning, energy load, rainfall, travel time, and inventory decisions all benefit from a useful estimate even when the estimate is imperfect.

A model can be valuable without being magical.

Science and Healthcare Need Context

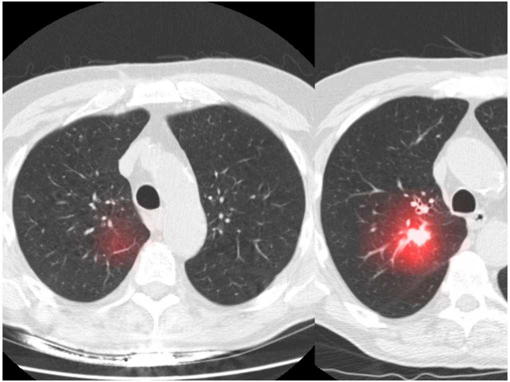

The same method meets different consequences

Medical support and scientific data analysis can both use AI, but the meaning of an error is very different from a bad movie recommendation.

Read this difference clearly

The important difference is not whether AI is present. It is what happens after a mistake: inconvenience, wasted money, delayed treatment, or a missed warning in research.

Why this matters early. Students should learn from lecture one that context is not an optional extra after the model is built.

Prediction Is Not Decision

Prediction is only one step

A model can estimate risk, but deciding what to do with that estimate is a separate step involving goals, constraints, cost, and responsibility.

A simple illustration

If a model predicts a high chance of late payment, the company still has to decide whether to send a reminder, change the payment plan, or involve a human reviewer.

Why this matters. Confusing prediction with decision makes AI look more autonomous than it really is.

Where Students May Meet AI

Not only in computer science

Students in science, health, business, design, education, and public policy all meet AI in different forms.

Possible starting points

Classifying images, finding unusual measurements, predicting a quantity, clustering survey responses, or supporting document analysis.

A good place to start

Start from a small, clear question where data exists and the meaning of success is easy to explain.

Why this is a better entry. A modest, well-scoped task teaches more than an ambitious idea with fuzzy boundaries.

Why Models Fail

Common reasons

Too little data, poor labels, a mismatch between training data and real use, overfitting, unstable evaluation, or the wrong task definition.

A better question to ask. If a model looks good in training and fails in use, the first question should not be "Which fancy model do we try next?" It should be "What exactly went wrong in the setup?"

How failure often looks

Models often fail quietly. They produce smooth outputs even when the reasoning behind them is weak.

High Stakes Change the Standard

When consequences rise, standards rise

The same error rate means different things in different settings. A wrong song recommendation is annoying. A wrong medical or driving judgment can be serious.

What changes in practice

Validation becomes stricter, human oversight matters more, and explanation matters more because the cost of a mistake is no longer trivial.

Why students should remember this. AI evaluation is inseparable from the cost of mistakes.

Good Numbers Can Mislead

Why metrics can fool us. An impressive accuracy score can still hide class imbalance, unfair behavior, weak robustness, or a test set that looks too much like the training set.

A classic trap

If only 1 in 100 cases is positive, a model can reach 99% accuracy by predicting “negative” almost every time and still be nearly useless.

Better habit. Ask which errors matter, what the baseline is, and whether the evaluation setting resembles the real use case.

Human Judgment Stays in the Loop

What AI does well

It can scan large pattern spaces, repeat routine judgments consistently, and surface cases worth attention.

What humans still add

Context, accountability, ethical judgment, exception handling, and the ability to ask whether the task itself was framed sensibly.

Useful sentence to keep. AI can support judgment without replacing responsibility.

Why This Course Starts Here

What we put on the table today

A working map: what AI means, why learning matters, what kinds of tasks appear most often, and what limits must be kept in view. **Why lecture 1 has two parts.** The AI half gives the outer map. The neural-network half gives one of the most important inner mechanisms.

What comes next

Data, perceptrons, multilayer networks, training, evaluation, large language models, and agents will all sit on top of the foundation from today. **A course reminder.** The point is not to worship the newest model. The point is to learn how to read a problem clearly.

Summary

- AI is a broad family of methods. Chatbots are only one visible piece of it.
- Machine learning became central because many real problems are too messy and too variable for hand-written rules.
- The hard part is not only building a model. It is asking the right question, choosing the right evidence, and judging mistakes in context.

Next Time

Where we go next. The next lecture moves from the outer map of AI to the inner logic of neural networks.

Keep this question in mind. If learning from data is so powerful, what kind of model can actually represent the patterns we care about?

Why that question matters. It leads directly to perceptrons, hidden layers, and the story of why neural networks returned to the center of modern AI.

Thank You