

# CSC 163: Computing, Ethics, and Society

## Chapter 2

### AI Fundamentals Exercises Solution

Department of Computer Science,  
Hofstra University

# Exercise: Task and learning type

---

For each scenario give the ML task (classification, regression, anomaly detection, sequence modeling; an agent acting in an environment = sequential decision making) and the learning type (supervised, unsupervised / self-supervised, reinforcement).

| Scenario                                                                          | Task | Learning type |
|-----------------------------------------------------------------------------------|------|---------------|
| Predict a house's sale price from its size and location, using past sales         |      |               |
| Label photos as 'cat' or 'dog' using human-labeled examples                       |      |               |
| Flag a network event that looks very unlike all normal traffic (no attack labels) |      |               |
| Predict the next word of a sentence, trained on raw text                          |      |               |
| A robot learns to walk through a maze by trial and error with rewards             |      |               |

# Exercise Solution: Task and learning type

---

Labeled data → supervised. No labels (patterns found in the data itself) → unsupervised / self-supervised. Predicting a number → regression; predicting a category → classification.

Next-word prediction is also a generative task: the model creates new data resembling its input.

| Scenario                                                                          | Task                       | Learning type                  |
|-----------------------------------------------------------------------------------|----------------------------|--------------------------------|
| Predict a house's sale price from its size and location, using past sales         | Regression                 | Supervised                     |
| Label photos as 'cat' or 'dog' using human-labeled examples                       | Classification             | Supervised                     |
| Flag a network event that looks very unlike all normal traffic (no attack labels) | Anomaly detection          | Unsupervised                   |
| Predict the next word of a sentence, trained on raw text                          | Sequence modeling          | Unsupervised / self-supervised |
| A robot learns to walk through a maze by trial and error with rewards             | Sequential decision making | Reinforcement                  |

## Exercise: Pipeline and model access

---

(a) The six steps of the ML pipeline are listed in scrambled order. Write 1 to 6 for the correct order.

(b) Name the access level for each case: closed, API or open.

(i) Users send queries through an interface with usage restrictions and cannot see the weights.

(ii) The weights are fully accessible and anyone may modify and redistribute them.

(iii) The model is used only inside the organization that built it.

| Step                  | Order (1–6) |
|-----------------------|-------------|
| Deployment            |             |
| Labeling              |             |
| Training              |             |
| Data collection       |             |
| Evaluation            |             |
| Choosing architecture |             |

# Exercise Solution: Pipeline and model access

---

- (a) Data collection → Labeling → Choosing architecture → Training → Evaluation → Deployment.
- (b) (i) **API** model; (ii) **open** model; (iii) **closed** model.
- Access level matters for safety: weights that are open cannot be taken back, while an API keeps access limited, monitored and revocable.

| Step                  | Order (1–6) |
|-----------------------|-------------|
| Deployment            | 6           |
| Labeling              | 2           |
| Training              | 4           |
| Data collection       | 1           |
| Evaluation            | 5           |
| Choosing architecture | 3           |

# Exercise: Types of AI

---

Match each description with one of: narrow AI, AGI, human-level AI (HLAI), transformative AI (TAI), superintelligence (ASI). (One type is defined by its impact on the world, not by its capabilities.)

| Description                                                                        | Type |
|------------------------------------------------------------------------------------|------|
| A chess engine that can do nothing except play chess                               |      |
| Performs many cognitive tasks across multiple domains at human or superhuman level |      |
| Capable of doing virtually everything humans do                                    |      |
| Systems whose impact on the world is at the level of the industrial revolution     |      |
| Outclasses human performance on virtually all intellectual tasks                   |      |

# Exercise Solution: Types of AI

---

TAI is the only type defined by impact (industrial-revolution scale); the others are defined by capability.

| Description                                                                        | Type                    |
|------------------------------------------------------------------------------------|-------------------------|
| A chess engine that can do nothing except play chess                               | Narrow AI               |
| Performs many cognitive tasks across multiple domains at human or superhuman level | AGI                     |
| Capable of doing virtually everything humans do                                    | Human-level AI (HLAI)   |
| Systems whose impact on the world is at the level of the industrial revolution     | Transformative AI (TAI) |
| Outclasses human performance on virtually all intellectual tasks                   | Superintelligence (ASI) |

# Exercise: True or false

---

For each statement write T or F. If it is false, correct it with a few words.

| Statement                                                                                                           | T / F | Correction (if false) |
|---------------------------------------------------------------------------------------------------------------------|-------|-----------------------|
| ML systems follow explicit pre-set instructions written for every decision                                          |       |                       |
| Supervised learning learns from unlabeled data                                                                      |       |                       |
| Unsupervised learning is also called self-supervised, because the model may generate supervisory signals for itself |       |                       |
| GPUs process sequentially and CPUs process in parallel                                                              |       |                       |
| Deep learning performance tends to taper off as more data is added                                                  |       |                       |
| Anomaly detection identifies outliers or abnormal data points                                                       |       |                       |

# Exercise Solution: True or false

---

Each statement: 1 point for T/F, 1 point for the correction.

| Statement                                                                                                           | T / F | Correction (if false)                                               |
|---------------------------------------------------------------------------------------------------------------------|-------|---------------------------------------------------------------------|
| ML systems follow explicit pre-set instructions written for every decision                                          | F     | They learn directly from data without explicit pre-set instructions |
| Supervised learning learns from unlabeled data                                                                      | F     | It learns from labeled data (labels are usually human-generated)    |
| Unsupervised learning is also called self-supervised, because the model may generate supervisory signals for itself | T     | —                                                                   |
| GPUs process sequentially and CPUs process in parallel                                                              | F     | GPUs: parallel processing; CPUs: sequential processing              |
| Deep learning performance tends to taper off as more data is added                                                  | F     | It tends to improve with more data (highly scalable)                |
| Anomaly detection identifies outliers or abnormal data points                                                       | T     | —                                                                   |