

Welcome to CSI 4106!

CSI 4106 - Fall 2026

Marcel Turcotte

Version: Sep 11, 2026 12:09

Preamble

Message of the day (MOTD)



Joëlle Pineau — an AI leader from Ottawa

- 2026-08 – [TIME100 AI 2026](#)

- 2026-06 – [uOttawa honorary doctorate](#)
- 2025-08 – Chief AI Officer, [Cohere](#)
- 2023-2025 – VP AI Research, Meta
- 2017-2023 – Led Facebook AI Research (FAIR), Montréal
- 2004-08 – [Professor](#), McGill University

[Cohere raises US\\$500-million, hires former Meta AI expert Joelle Pineau](#), Joe Castaldo, The Globe and Mail, 2025-08-14.

[Joëlle Pineau](#) is a Canadian AI researcher born in Ottawa. A woman who has become an international leader in AI, she offers students a local example of the people shaping this field. She is Chief AI Officer at Cohere, a professor at McGill University and a core academic member of [Mila](#). Previously, she led Meta’s Fundamental AI Research (FAIR) team as Vice President of AI Research.

Her work includes reinforcement learning and robotics: how can a machine learn to make decisions through interaction with its environment? Applications include health care and assistive robots. She also champions reproducibility: reporting methods, code and experimental details so others can check and build on research findings. These are useful connections to the algorithms and experimental practices we will study in this course.

She received an honorary doctorate from the University of Ottawa at the June 2026 convocation, recognizing her contributions to AI research and responsible, open science. She was also named to TIME100 AI in August 2026.

[Cohere](#) is a Canadian AI company that builds language models, search tools and AI assistants for businesses and public-sector organizations. Its products include Command language models and North, a workplace AI platform. For example, an organization can use these tools to search its internal documents and answer questions about them. Cohere emphasizes privacy, security and deployment options that let organizations retain control of their data and infrastructure. The company recently acquired the German firm Aleph Alpha to accelerate its European expansion in the field of artificial intelligence consulting services, as part of a strategic transaction backed by an \$800 million investment ([Canadian AI firm Cohere, Germany’s Aleph Alpha announce merger](#)).

Cohere’s cofounders are Aidan Gomez, Nick Frosst and Ivan Zhang. Gomez co-authored the 2017 paper “[Attention Is All You Need](#),” which introduced the transformer architecture underlying many modern language models.

Sources: [uOttawa biography](#), [2026 honorary doctorate recipients](#), [NSERC research profile](#), [McGill reproducibility checklist](#), and [Ottawa birthplace profile](#).

- [Federal government taps Cohere to work on use of AI in public service](#), Anja Karadeglija, The Canadian Press, 2025-08-20.
- [Ceremony 13 - Faculty of Engineering](#): Watch Joëlle Pineau's convocation address on June 12, 2026 (starts at 1 h 23 m 24 s).

<https://youtu.be/Sq1QZB5baNw>

The presentation includes four videos, but we will only watch one. These videos were chosen for their science fiction aspect and their use of generative AI.

Figure's approach is markedly different from earlier methods. The latter often relied on separate modules to solve different problems. For example, a robot would use a planning algorithm to find an optimal path, and the reasoning was heavily scripted ("on rails"), thus limiting tasks and learning capabilities.

Here, we see an end-to-end approach based on neural networks. Deep learning aims to produce flexible, general, and capable learning systems. Inspired by large language models (LLMs), we are now observing the development of large action models (LAMs).

The videos on the new models from Google DeepMind are equally impressive.

How do you feel about these developments?

What is your favorite AI video? Please share the information in the general forum on Brightspace or email it to me.

<https://youtu.be/Z3yQHYNXPws>

<https://youtu.be/UIZAiXYceBI>

<https://youtu.be/nXVvvRhiGjl>

Learning objectives

- **Clarify** the proposition
- **Discuss** the syllabus
- **Articulate** the expectations
- **Explore** the various definitions of "artificial intelligence"

I want to clarify my proposal. I have chosen a specific approach to introduce the concepts and I would like to explain the reasons for this choice.

After presenting the course outline and expectations, we will discuss the different definitions of artificial intelligence.

Proposition

Course overview

Calendar description

The roots and scope of Artificial Intelligence. Knowledge and knowledge representation. Search, informed search, adversarial search. Deduction and reasoning. Uncertainty in Artificial Intelligence. Introduction to Natural Language Processing. Elements of planning. Basics of Machine Learning.

Above is the official course description, which dates back several years, placing machine learning at the end of the list.

At the end of this presentation, you will find the Python code I used to produce this audio clip.

Aims: Deep learning early

(Wang 2019)

To the larger community of computer science and information technology, AI is usually **identified by the techniques grown from it**, which at **different periods** may include **theorem proving, heuristic search, game playing, expert systems, neural networks, Bayesian networks, data mining, agents**, and recently, **deep learning**.

- **Deep learning** is so dominant that I have chosen to structure everything around it

What does this mean?

Good Old-Fashioned AI (GOFAI) relied on **hand-crafted knowledge engineering**, but it has been largely displaced by **machine learning** due to the increased availability of data, computing resources, and new algorithms.

...

Deep learning has significantly impacted various domains, including natural language processing, robotics, and computer vision.

...

However, deep learning has current **limitations**, particularly in **reasoning**, where symbolic AI excels and could potentially offer valuable insights.

But also

EVOLUTION, AI, AND THE
FIVE BREAKTHROUGHS
THAT MADE OUR BRAINS

**A BRIEF HISTORY
OF INTELLIGENCE**

MAX BENNETT

- In [A Brief History of Intelligence](#) (Bennett 2023), Max Bennett discusses **significant milestones in the evolution of human intelligence** and draws parallels to advancements in artificial intelligence (AI).
- Learning itself represents one of the earliest and most extensively understood milestones in the evolution of intelligence.

Curiously, the development of AI has been largely influenced by logical approaches (symbolic AI). The development of AI has been strongly marked by intellectual currents rooted in philosophy and mathematics, rather than in biology and its evolution. Reasoning in philosophy and mathematics relies on complex cognitive functions, perhaps less well understood and having evolved later.

Learning itself represents one of the first and most widely understood steps in the evolution of intelligence.

It would have been logical to approach the study of intelligence by progressing from simpler forms to more complex ones.

Aims: Applied

Andrew Ng, [The Batch, April 10, 2024](#)

Many software developers worry that **large language models** will make **human coders obsolete**. We doubt that AI will replace coders, but we believe that **coders who use AI will replace those who don't**.

- Whenever possible, concepts will be introduced with code.

Aims: Applied

Richard Feynman

What I cannot build. I do not understand.

Course Philosophy and Expectations

- **Theory to Practice**
- The Necessity of **Code** and **Mathematics**
 - **Note:** data visualization code is exempt from examinations.
- **Active** Engagement
- The Value of **Inquiry**

Course Philosophy and Expectations

- **Theory to Practice:** To build a rigorous understanding of fundamental concepts, we will implement core algorithms from the ground up. For example, you will construct logistic regression to master gradient descent and implement backpropagation to understand the mechanics of deep learning. Concurrently, we will leverage industry-standard libraries, such as **scikit-learn**, **Keras**, and **TensorFlow**, to address complex, real-world problems.

- **The Necessity of Code and Mathematics:** A deep understanding of artificial intelligence requires active engagement with both the underlying mathematics and the Python implementations. Skipping these foundational elements will leave substantial gaps in your knowledge. Experience dictates a clear pattern: students who actively read, write, and execute code achieve a superior, lasting comprehension of the material.
 - **Note:** While understanding the core algorithmic code is essential, data visualization code is provided strictly for pedagogical support and is exempt from examinations.
- **Active Engagement:** Mastery results from the combination of structured course design and your rigorous effort. You are expected to thoroughly review all materials, execute the provided computational examples, attend lectures, and participate actively.
- **The Value of Inquiry:** A large cohort brings a diversity of perspectives that enriches the academic environment. Bring your questions to class; your feedback is continuously integrated into the lecture notes and frequently informs the development of exam questions.
- **Future Outcomes:** This dual approach of theoretical grounding and practical application is proven. Alumni consistently report that this methodology effectively prepared them for both industry careers and advanced graduate research.

Web sites

- turcotte.xyz/teaching/csi-4106
 - [Schedule](#)
 - [Syllabus](#)
 - mirror: turcotte.github.io/csi-4106
- uottawa.brightspace.ca

Course information is provided on two websites: Brightspace and my personal website. The personal website includes the syllabus, lecture notes (slides), useful links, and a FAQ. To ensure continuous availability, the same content is mirrored on github.io at turcotte.github.io/csi-4106. Brightspace is used for assignment submission and participation in discussion groups.

Beta testers

This will be my third iteration of this content. Your help identifying what works and what doesn't will be most appreciated.

Warnings

CSI 4106 is an introductory course on artificial intelligence, offering a brief overview of various topics within this broad field. **Each topic covered could be explored in much greater depth through one or more graduate-level courses.** The primary objective of CSI 4106 is to provide students with a foundational understanding of the core areas that constitute artificial intelligence.

...

Overlaps with other courses are inevitable, but I will do my best to keep it at a minimum.

...

This is not a course on the impact of AI on society, including ethics, fairness, trust and safety.

This warning is actually for myself. These are topics I am passionate about, and I would love to share everything I know with you. However, that is obviously not possible.


Generally, I will try to focus on a small number of approaches to thoroughly understand the subjects, rather than adopting an exhaustive approach.

The University of Ottawa offers courses on the ethics of artificial intelligence. I present here two that I am familiar with.

- [CSI 5195 Ethics for Artificial Intelligence](#)
- [IAI 5130 Ethics for Design, AI and Robotics](#)

Setting the Stage: AI, Deep Learning, and Diverging Views on Intelligence.

AI, ML, DL



Attribution: Lily Popova Zhuhadar, [CC BY-SA 4.0](#), via Wikimedia Commons, 2026-09-08.

Deep learning is so prevalent today that some people might confuse it with artificial intelligence. As the figure shows, deep learning is one of many techniques used in machine learning. Machine learning, in turn, is one of several disciplines within artificial intelligence. Other AI disciplines include knowledge representation, reasoning and planning, natural language processing, computer vision, and robotics.

By the end of this course, this distinction should be very clear.

Influential Approaches

- **Symbolic AI:** represents knowledge using explicit **symbols** and **structures**, manipulated through **rules**, **inference**, or **search**.
- **Connectionist AI:** represents knowledge through **patterns of activity** and **connection weights in networks** of computational units.

These approaches can be combined, and they do not encompass all of AI.

At the outset of this course, it is important to recognize that two main approaches exist in AI: **symbolic AI** and **connectionism**. Initially, the symbolic approach was dominant in the field of AI, but today, the connectionist approach prevails.

Any attempt to impose a rigid taxonomy on complex concepts is likely to obscure as much as it reveals. In *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (Domingos 2018), Pedro Domingos identifies five “tribes” of machine learning: symbolists, connectionists, evolutionaries, Bayesians, and analogizers. Inevitably, some concepts and algorithms do not fit neatly into any single category. Moreover, some approaches are intentionally designed as hybrid methods.

Symbolic systems in mathematics

- **Algebra**
 - $x + 3 = 7$, here x represents a quantity, x is a symbol.
- **Calculus**
 - Symbolic differentiation of x^2 into $2x$, x is a symbol.
- **Logic**
 - We might use P to mean “It is raining” and Q mean “The ground is wet.” P and Q are symbols.

You already work with symbolic systems in mathematics. In algebra, x represents a quantity, and rules let you transform expressions such as $x + 3 = 7$ into $x = 4$. In symbolic AI, symbols can also represent objects and relationships. For example, `Parent(alan, chris)` represents a relationship, and a rule can infer a grandparent relationship from two parent relationships.

Modus ponens (from logic) gives a simple example of symbols and a rule of inference:

Let P mean “It is raining” and Q to mean “The ground is wet.”

$P \rightarrow Q$ If it is raining, the ground is wet.

P It is raining.

$\therefore Q$ Therefore, the ground is wet.

Here, P and Q **represent propositions**, and $\rightarrow$ **expresses implication**. The inference rule operates on this explicit structure: from $P \rightarrow Q$ and P , infer Q .

We also use algebra and calculus to design neural networks. What makes an approach symbolic is its explicit representation and manipulation of meaningful expressions, not simply its use of mathematical notation.

Towers of Hanoi

(for your information only)

<https://youtu.be/PGuRmqpr6Oo>

See also: Binary, Hanoi and Sierpinski, [Part 1](#) and [Part 2](#), by [3Blue1Brown](#).

Symbolic AI (Planning)

Problem

The **Towers of Hanoi** is a puzzle that consists of **three pegs** and **a number of disks of different sizes**. The puzzle **starts** with all the **disks stacked in decreasing size on one peg**, and the **goal is to move the entire stack to another peg**, following these rules:

1. **Only one disk** can be moved at a time.
2. A disk can only be placed **on top of a larger disk** or **on an empty peg**.

Rules

Action `Move(X,Y,Z)`:

Preconditions = {`Clear(X)`, `On(X,Y)`, `Clear(Z)`, `Smaller(X,Z)`};

Effects = {`-On(X,Y)`, `Clear(Y)`, `On(X,Z)`, `-Clear(Z)`};

Start

`D1`, `D2`, `D3`, `P1`, `P2`, `P3` are **symbols**, where `D1`, `D2`, and `D3` are disks, and `P1`, `P2`, and `P3` are pegs.

```
On(D1, D2), On(D2, D3), On(D3, P1),  
Clear(D1), Clear(P2), Clear(P3),  
Smaller(D1, D2), Smaller(D1, D3), Smaller(D2, D3),  
Smaller(D1, P1), Smaller(D1, P2), Smaller(D1, P3),  
Smaller(D2, P1), Smaller(D2, P2), Smaller(D2, P3),  
Smaller(D3, P1), Smaller(D3, P2), Smaller(D3, P3).
```

Goal

```
On(D1, D2), On(D2, D3), On(D3, P3).
```

Solution

```
Move(D1, D2, P3)  
Move(D2, D3, P2)  
Move(D1, P3, D2)  
Move(D3, P1, P3)  
Move(D1, D2, P1)  
Move(D2, P2, D3)  
Move(D1, P1, D2)
```

To complete the code, it should be specified that a peg is considered to be wider than all of the disks combined.

STRIPS (Stanford Research Institute Problem Solver, 1971) is a **classical AI planning system** that represents problems in terms of:

- **States:** sets of logical predicates describing the world.
- **Operators (Actions):** defined by
 - *Preconditions:* predicates that must be true to apply the action,
 - *Add list:* predicates made true after execution,
 - *Delete list:* predicates made false after execution.
- **Goal:** a set of predicates to be achieved.

A plan is found by applying actions to transform the initial state into a goal state, updating the world via the add/delete lists.

These methods will be examined in depth starting from Lecture 16, focused on search methods.

Symbolic AI (Kinship)

Problem

Given a few **facts** about who is parent of whom (**symbols**) and a handful of **Horn-clause rules**, infer new relations (e.g., **ancestor**, **siblings**, **grandparent**) by *logical deduction*.

Rules

```
father(F,C) :- male(F), parent(F,C).
mother(M,C) :- female(M), parent(M,C).

sibling(X,Y) :- parent(P,X), parent(P,Y), X \= Y.

grandparent(G,C) :- parent(G,P), parent(P,C).

ancestor(A,D) :- parent(A,D).
ancestor(A,D) :- parent(A,X), ancestor(X,D).
```

Facts

```
% facts/symbols
```

```
male(alan).    female(brenda).
male(chris).   female(dina).
male(eli).     female(fiona).

parent(alan, chris).
parent(brenda, chris).
parent(chris, dina).
```

```
parent(dina, eli).
parent(dina, fiona).
```

Query

```
?- grandparent(G, dina).
% Expected: G = alan ; G = brenda.
```

```
?- ancestor(A, fiona).
% Expected: A = dina ; A = alan ; A = brenda ; A = chris ;
```

```
?- sibling(eli, fiona).
% Expected: true.
```

The preceding example demonstrates the application of logic programming, specifically Prolog, to articulate relationships between entities.

- **Clear separation** of knowledge (facts) and reasoning (rules).
- **Explanations:** Prolog's proof trees make deductions inspectable.
- **Recursion** illustrates expressive power (e.g., ancestor/2).

Symbolic AI (Prerequisites)

Problem

Given **course prerequisites** and a student's **completed** set, infer which courses they **can take next**. This shows *symbolic constraint reasoning*.

Rules

```
all_prereqs_met(Student, Course) :-
    \+ (prereq(Course, P), \+ completed(Student, P)).
```

```
can_take(Student, Course) :-
    all_prereqs_met(Student, Course),
    \+ completed(Student, Course).
```

Courses

```
% --- Course graph (facts, symbols) ---
```

```
prereq(csi2120, csi2110).
prereq(csi2110, iti1121).
prereq(csi2110, mat1338).
prereq(iti1121, iti1120).
```

Students

```
% --- Student record (facts, symbols) ---
```

```
completed(alex, iti1121).  
completed(alex, iti1120).  
completed(alex, mat1338).
```

Query

```
?- can_take(alex, csi2110).  
% true
```

```
?- can_take(alex, csi2120).  
% false
```

```
?- can_take(alex, iti1121).  
% false
```

```
?- all_prereqs_met(alex, csi2110).  
% true
```

Another example using logic programming (Prolog).

- Shows **constraint satisfaction** and **closed-world negation**.
- Easy to extend with electives, co-requisites, or anti-requisites.
- Produces **auditable recommendations** (why/why not).

In Prolog, negation as failure is represented as `\+` (backslash-plus).

- `\+ Goal` succeeds if `Goal` cannot be proven.
- Semantics: **closed-world assumption** — if it's not in the knowledge base, we assume it's false.

The rule

```
{prolog}  
all_prereqs_met(Student, Course) :-  
    \+ (prereq(Course, P), \+ completed(Student, P)).
```

reads as:

"all_prereqs_met(Student, Course) holds **if there does not exist** a prerequisite `P` of `Course` such that `Student` has not completed `P`."

Or, more naturally:

"A student has met all the prerequisites for a course when every course that is listed as a prerequisite has already been completed by that student."

Symbolic AI: historical perspective

- “Their founding tenet held that **knowledge** can be represented by a **set of rules**, and computer programs can use **logic** to manipulate that knowledge.” (Strickland 2021)
- “Researchers developing symbolic AI set out to **explicitly teach computers** about the world.” (Strickland 2021)
- “(. . .) a **physical symbol system** has the **necessary** and **sufficient** means for **general intelligent action**.”
(Newell and Herbert A. Simon 1976b)

In these early systems, people often supplied knowledge explicitly as facts and rules. However, symbolic representations can also be learned from examples. The defining distinction is the use of explicit symbols and operations over them, rather than whether learning occurs.

The researchers of the time were convinced that the symbolic approach was the key to success.

Symbolic AI

- What do you think made **symbolic AI difficult** in practice?

The main challenges of symbolic AI are usually framed as follows:

- **Knowledge acquisition bottleneck**
 - It is hard to manually encode all the facts, rules, and exceptions a system needs (Feigenbaum 1977).
- **Brittleness**
 - Symbolic systems often work well only in narrowly defined domains and fail when conditions shift or inputs are noisy/ambiguous (Lenat and Guha 1989)
- **Combinatorial explosion**
 - Search and inference can become intractable as the number of symbols and rules grows (Lighthill 1973).
- **Handling uncertainty**
 - Classical symbolic AI is not naturally good at probabilistic, incomplete, or vague information (Pearl 1988).
- **Perception and grounding**
 - Symbols can be detached from the physical world or sensory input, making it hard to connect abstract representations to real-world meaning (Harnad 1990).
- **Commonsense reasoning**
 - Humans rely on vast amounts of implicit commonsense knowledge that is difficult to formalize explicitly (McCarthy and Hayes 1969).
- **Acquiring knowledge:**
 - Manually engineered systems require people to supply and maintain their knowledge. Symbolic learning methods can acquire rules from examples,

although choosing suitable representations remains a challenge.

- **Context sensitivity**
 - Human understanding depends heavily on context, but symbolic rules are often too rigid to capture this (Dreyfus 1972).

However, the researchers of the time were convinced that the symbolic approach was the key to success (Newell and Herbert A. Simon 1976a).

Connectionist

Artificial neural networks (ANNs) are computational models loosely inspired by biological neural systems. They consist of interconnected computational units with weighted connections.

...

During **training**, their parameters, including connection weights, are adjusted to improve performance on a specified objective.

...

Knowledge is encoded in the **network's parameters** and **patterns of activity**, rather than primarily as explicit facts and rules.

See: playground.tensorflow.org


The term "connectionists" comes from the idea that nodes in these models are interconnected. Instead of being explicitly programmed, these models learn their behavior through training.

Designers still choose the architecture, training objective, and learning procedure. Training adjusts parameters rather than requiring designers to specify every input–output response. These models are not detailed simulations of the human brain.

Deep learning is a connectionist approach.

A hybrid system: A neural network identifies objects in an image; a symbolic planner uses those objects and explicit action rules to construct a plan.

Connectionist



Attribution: LeNail, (2019). NN-SVG: Publication-Ready Neural Network Architecture Schematics. Journal of Open Source Software, 4(33), 747, <https://doi.org/10.21105/joss.00747> (GitHub)

Definig AI

Survey

Perceptions and Attitudes Toward Artificial Intelligence.

- The majority of students exhibit a generally positive disposition towards AI. If your perspective differs, please elaborate.
- There is a wide range of opinions regarding the probability of catastrophic outcomes due to AI (P(doom)).
- A relatively small number of students express concern about AI replacing their jobs.
- Few students perceive current AI systems as genuinely intelligent.
- Views on Artificial General Intelligence (AGI) and Super-intelligence vary significantly.
- The primary concerns surrounding AI include potential job displacement and misuse, with privacy issues also noted.
- Many students regard the medical and healthcare sectors as areas where AI could exert significant influence.
- A substantial portion of students believes that AI is excessively hyped.
- Fewer than 20% of the students possess substantial experience in artificial intelligence.

Why defining AI?

- **Enhancing Clarity and Communication:** Establishing a mutual understanding of evaluation criteria, benchmarking tests, and milestones.
- **Addressing Ethical, Social, and Legal Implications:** Enabling policymakers to develop regulations that minimize AI-related risks.

An operational definition

(Wang 2019)

To the larger community of computer science and information technology, AI is usually identified by the techniques grown from it, which at different periods may include **theorem proving, heuristic search, game playing, expert systems, neural networks, Bayesian networks, data mining, agents**, and recently, **deep learning**.

Context

- Significant **hype**.
- Several leading companies, including [DeepMind](#) and [OpenAI](#), are primarily focused on the development of **artificial general intelligence** (AGI).
- Should AGI become a reality, **its implications could profoundly impact various aspects of our lives**, including education, employment, economic structures, and scientific advancements.

AI thus has great potential to influence the course of your lives. It is therefore relevant to take the time to define what AI is and to explore its different approaches.

I should warn you right away that there is no consensus on the definition of AI, nor on the definition of intelligence.

Daily **Mail**
.com

21 June 2024

**ChatGPT passes the famous 'Turing test'
- suggesting the AI bot has intelligence
equivalent to a human, scientists claim**

Psychology Today

May 17, 2024

GPT-4 passes the Turing test, marking a potential inflection point in AI's mastery of human-like language.

Stanford School of
Humanities and Sciences

FEBRUARY 22, 2024

The most recent version of ChatGPT passes a rigorous Turing test

Attribution: Melanie Mitchell, Substack post titled: [The Turing Test and Our Shifting Conceptions of Intelligence](#), August 15, 2024.

4 days ago

Question

User: The man and the goat want to enjoy a picnic near a river, but there's a wolf. What to do with the zucchini?

...

ChatGPT 5: You've got the classic river-crossing (wolf-goat-cabbage ⇒ here, zucchini).

Boat carries you + one item. Unsafe pairings: wolf with goat; goat with zucchini.

Minimal plan:

1. Take goat over.
2. Return alone.
3. Take wolf over.
4. Bring goat back.
5. Take zucchini over.

6. Return alone.
7. Take goat over.

Conclusion: keep the zucchini never alone with the goat; ferry it on step 5.

...

Attribution: Andriy Burkov, published on [LinkedIn](#) August 2025.

Question

User: Surgeon but child, hopsital, it have a no, why it so?

...

ChatGPT 5: You're pointing to the classic riddle:

Car crash → injured child → surgeon says "I can't operate—this is my child."

Resolution: the surgeon is the child's mother (or another parent in a non-heteronormative family).

The "paradox" arises from implicit gender assumptions about surgeons.

...

Attribution: Andriy Burkov, published on [LinkedIn](#) 3 weeks ago.

These two examples illustrate that large language models generate output by leveraging statistical associations between words (tokens).

The complex nature of intelligence

...

- **"Artificial"** pertains to the creation of entities or phenomena that mimic natural processes using technology or synthetic materials, a definition broadly recognized and accepted.
- Therefore, defining **"artificial intelligence"** fundamentally requires us to first clarify what we mean by **"intelligence."** Surprisingly, "[d]espite a long history of research and debate, there is still no standard definition of intelligence." (Legg and Hutter 2007)

How do you define intelligence?

- What are the **characteristics** you associate with **intelligence**?
- “**Intelligence** is a very general mental capability that, among other things, involves the ability to **reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience.**”
 - Originally published in the Wall Street Journal in 1994, the content was reprinted in 1997: (Gottfredson 1997)

Can AI make mistakes?

Thinking, acting, humanly, rationally

Russell and Norvig considers two axes: thinking vs behaviour, human vs rationality.

	Thinking	Acting
Human-like	thinking humanly (simulation)	acting humanly (Turing test)
Rationality	thinking rationally (logic)	acting rationally (agent)

See also the appendix – [Section 8](#) On Defining Artificial Intelligence

The question of intelligence has been the subject of much debate in the literature and the media, particularly when it comes to animals or computers. This abundance of information can bias our thinking.

A simple thought experiment might offer a new perspective: would you be able to recognize intelligence in an extraterrestrial entity?

Rationality




(Mohammed et al. 2019)

Rationality involves the **evaluation of choices to achieve a goal or to find the optimal solution to a problem**. Simon (1972, p. 161) defined rationality as "a style of behavior that is appropriate to the **achievement of given goals**, within the **limits imposed by given conditions and constraints**."

Attribution: [NBC Television](#), [Public domain](#), via [Wikimedia Commons](#)

Cognitive Taxonomy



Deconstructing general intelligence

Cognitive Taxonomy (continued)

- **Perception:**
 - The ability to extract and process **sensory information** from the environment.
- **Generation:**
 - The ability to produce outputs such as **speech, text, motor movements**, and computer control **actions**.
- **Attention:**
 - The ability to **focus** cognitive resources on specific aspects of perceptual stimuli, thoughts, or task demands.

Burnell et al. (2026)

Cognitive Taxonomy (continued)

- **Learning:**
 - The ability to **acquire** new **knowledge, skills**, or **understanding** through **experience, study**, or **instruction**.
- **Memory:**
 - The ability to **store** and **retrieve** information over time.
- **Reasoning:**
 - The ability to **draw valid conclusions** and **make inferences** by applying logical principles.

Cognitive Taxonomy (continued)

- **Metacognition:**
 - The **knowledge** a system has about its **own cognitive processes** and its ability to **monitor** and **control** those processes.
- **Executive functions:**
 - Abilities that facilitate **goal-directed behavior**. Includes **planning, inhibition,** and cognitive flexibility.

Cognitive Taxonomy (continued)

- **Problem solving:**
 - The ability to **find effective solutions** to **domain-specific problems**.
- **Social cognition:**
 - The ability to **process** and **interpret** social information and to **respond appropriately** in social situations.

For Burnell et al. (2026), problem-solving and social cognition require the integrated application of the faculties described above.

Here are some of the views commonly held by prominent researchers in the field.

Douglas Hofstadter

No one knows where the borderline between non-intelligent behavior and intelligent behavior lies, in fact, to suggest that a sharp border exists is probably silly. But essential abilities for intelligence are certainly:

1. to respond to situations very **flexibly**;
2. to take advantage of **fortuitous** circumstances;
3. to make sense out of **ambiguous** or **contradictory messages**;
4. to recognise the **relative importance of different elements** of a situation;
5. to find **similarities** between situations despite differences which may separate them;
6. to draw **distinctions** between situations despite similarities which may link them;
7. to **synthesize** new concepts by taking old concepts and putting them together in new ways;
8. to come up with **ideas which are novel.**"

For certain complex concepts, drawing a clear boundary can prove challenging. Take the concept of life, for example. Humans, plants, and insects are considered living, as are microorganisms such as bacteria. However, viruses and viroids are not.

François Chollet, Creator of Keras

Real intelligence is not about mastering an individual skill, he argued, but about **taking what has been learned and applying it to a new, different situation.**

In his view, intelligence is the ability to **efficiently acquire new skills** that training did not prepare for, with the goal of accomplishing tasks that are sufficiently different from those a system has seen before.

The wider the scope of the new skills, the closer the computer comes to achieving artificial general intelligence.

"If you can make the learning process as information-efficient as a human mind, then you've got AGI," Chollet said.

So far, machines lag far behind, approximately 10,000 times less efficient than human brains. For instance, **it took millions of images to teach computers to recognize pictures of cats, whereas humans learn to identify them based on only one or two examples.**

Savage (2024)

Narrow vs General AI

OpenAI GPT-6 Astra

[OpenAI GPT-6 Astra](#) achieves SOTA on [ARC-AGI](#)

Artificial General Intelligence (AGI)

Artificial general intelligence (AGI) refers to a form of artificial intelligence (AI) that either **equals** or **exceeds human proficiency** across a **diverse array of cognitive functions.**

AKA **human-level intelligence.** As opposed to **narrow intelligence**, the current status of AI, which is designed to perform a specific task or a limited range of tasks.

AlphaFold (1, 2, & 3)

I repeat, **there is nothing wrong with narrow AI.**

- «Two papers in this week's issue dramatically expand our structural understanding of proteins. Researchers at **DeepMind**, Google's London-based sister company, present the latest version of their **AlphaFold** neural network.»

- Jumper et al. (2021)

In narrow AI, several systems already surpass human performance in specific domains, including [AlphaZero](#), [Stockfish](#), and [AlphaFold](#).

Accurate weather prediction is crucial because it provides advanced warning of severe events, enabling timely evacuations and reducing casualties and damage. It also supports key sectors such as agriculture, aviation, shipping, logistics, and energy, while helping manage public health risks like disease spread, poor air quality, and heat-related crises. Physics-based models are computationally intensive, limiting forecast generation frequency. In recent years, AI-based weather models have been developed that achieve strong medium-range forecasting performance while offering substantial computational efficiency. Google DeepMind recently released WeatherNext3, which demonstrates improved spatial and temporal coverage. WeatherNext3 is an example of narrow artificial intelligence.

- [Introducing WeatherNext 3, our most advanced and accurate global weather AI model](#) (blog)
- [WeatherNext 3: More accurate, timely, and local weather forecasts](#) (video)
- [WeatherNext 3: Increasing resolution and performance of global weather models with raw observations](#) (paper)

AI effect/paradox

(Wang 2019)

(. . .) as soon as a computer system is built to solve a problem successfully, the problem is no longer "only solvable by the human mind," so does not need intelligence anymore. Consequently, "**AI is whatever hasn't been done yet**" (Hofstadter, 1979; Schank, 1991), which is known as "**the AI Effect**" (McCorduck 2004).

The paradox of AI is quite fascinating. In the early days of AI, researchers focused on problems like differential equations or chess.

Each time a computer solves one of these major problems, we come to think that perhaps that problem didn't really require intelligence to be solved after all.

The defeat of Garry Kasparov at the hands of IBM's Deep Blue is a perfect example.

Impact

Economy

Beyond the hype: Capturing the potential of AI and gen AI in tech, media, and telecom 2024-02-22

McKinsey research estimates that **gen AI** could add to the economy between *2.6trillion* and **4.4 trillion annually** while increasing the impact of all artificial intelligence by 15 to 40 percent.

In fact, it seems possible that **within the next three years, anything not connected to AI will be considered obsolete or ineffective.**

Subfields of AI

1. **Machine Learning:** Credit card fraud detection
2. **Deep Learning:** Image and facial recognition
3. **Natural Language Processing:** Virtual assistants like Siri or Alexa
4. **Computer Vision:** Autonomous vehicles
5. **Robotics:** Industrial automation in manufacturing
6. **Expert Systems:** Medical diagnosis support
7. **Speech Recognition:** Voice-to-text transcription services
8. **Planning and Decision Making:** Supply chain optimization
9. **Reinforcement Learning:** Game AI in complex strategy games
10. **Knowledge Representation:** Semantic web technologies for information retrieval

Our Final Invention

(Russell and Norvig 2020)

AI expert Kai-Fu Lee predicts that its **impact** will be **“more than anything in the history of mankind.”**

<https://youtu.be/ixgunKpy61s>

What Does the AI Boom Really Mean for Humanity? | The Future With Hannah Fry.
Bloomberg Originals, posted on YouTube on 2024-09-12.

https://youtu.be/sK5_pQV_QEA

Questions

- Can the concept of intelligence be considered **independently** of the entities that express it? This is the problem of **embodiment**.
- Can a **machine** exhibit human-level intelligence?
- Is it possible to **dissociate** the following concepts from that of intelligence?
 - **Agency.**
 - **Sentience.**
 - **Consciousness.**
 - **Emotions.**
 - **Language.**
 - **Mind.**
- Can an AI **suffer**?

For me, an interesting definition of intelligence would be similar in nature to that of computation.

Using **computation** as an analogous concept can help clarify this question. Theoretically, computation is often considered in an abstract manner, independent of any physical implementation. The Church–Turing thesis states that any effectively calculable function can be computed by a Turing machine. This provides a way to discuss computation abstractly, without committing to a particular physical implementation.

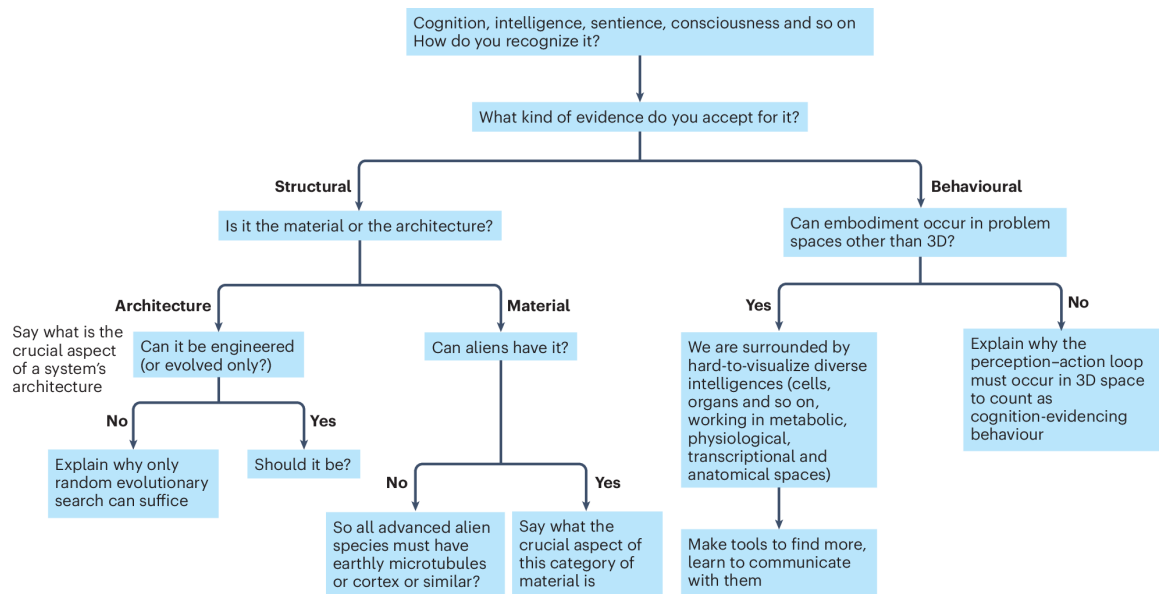
I must confess: until 2022, I was rather skeptical. I often drew a parallel between artificial intelligence and alchemy. Alchemists never succeeded in transmuting lead into gold because changing lead into gold requires nuclear transformation, not ordinary chemical reactions, but their approach led to the development of chemistry. Similarly, we may never be able to replicate human intelligence, but the impacts of AI are immense and often unexpected, like the garbage collector introduced by Lisp.

Then, it struck me: I am a scientist, and I wondered what it would mean if we were not able to produce intelligence comparable to that of humans.

Richard Feynman, who received the Nobel Prize in Physics in 1965, famously said, “What I cannot build. I do not understand.”

Why seek to dissociate intelligence from other concepts such as agency, consciousness, and emotions? A computer system devoid of these attributes seems, indeed, far less daunting.

Deepen the Reflection



Rouleau, N. & Levin, M. (2024). Discussions of machine versus living intelligence need more clarity. *Nature Machine Intelligence*, 6(12), 1424–1426.

Prologue

Summary

- **Discussed** the syllabus
- **Distinguish** the concept of artificial intelligence from the concept of machine learning
- **Distinguish** symbolic AI from connectionist AI
- **Explored** the various definitions of "artificial intelligence"

Next lecture

- Introduction to machine learning

References

Bennett, Max S. 2023. *A Brief History of Intelligence: Evolution, AI, and the Five Breakthroughs That Made Our Brains*. First edition. Mariner Books.

Burnell, Ryan, Yumeya Yamamori, Orhan Firat, et al. 2026. "Measuring Progress Toward AGI: A Cognitive Framework." *arXiv*, ahead of print.

<https://doi.org/10.48550/arxiv.2605.28405>.

Domingos, Pedro. 2018. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. Basic Books, Inc.

Dreyfus, Hubert L. 1972. *What Computers Can't Do: The Limits of Artificial Intelligence*. Harper & Row.

Feigenbaum, Edward A. 1977. "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering." *Proceedings of the 5th International Joint Conference on Artificial Intelligence* 2: 1014–29.

Gottfredson, Linda S. 1997. "Mainstream Science on Intelligence: An Editorial with 52 Signatories, History, and Bibliography." *Intelligence* 24 (1): 13–23.
[https://doi.org/10.1016/s0160-2896\(97\)90011-8](https://doi.org/10.1016/s0160-2896(97)90011-8).

Harnad, Stevan. 1990. "The Symbol Grounding Problem." *Physica D: Nonlinear Phenomena* 42 (1-3): 335–46.

Jumper, John, Richard Evans, Alexander Pritzel, et al. 2021. "Highly accurate protein structure prediction with AlphaFold." *Nature*, 1–11. <https://doi.org/10.1038/s41586-021-03819-2>.

Legg, Shane, and Marcus Hutter. 2007. "A Collection of Definitions of Intelligence." *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms: (NLD)*, 17–24. <https://doi.org/10.5555/1565455.1565458>.

Lenat, Douglas B, and Ramanathan V Guha. 1989. *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Addison-Wesley.

Lighthill, James. 1973. *Artificial Intelligence: A General Survey*. Science Research Council.

McCarthy, John, and Patrick J Hayes. 1969. "Some Philosophical Problems from the Standpoint of Artificial Intelligence." In *Machine Intelligence 4*, edited by Bernard Meltzer and Donald Michie. Edinburgh University Press.

McCorduck, Pamela. 2004. *Machines Who Think, A Personal Inquiry into the History and Prospects of Artificial Intelligence*. Taylor & Francis Group, LLC.
<https://doi.org/10.1201/9780429258985>.

Mohammed, Anne-Marie, Sandra Sookram, and George Saridakis. 2019. "Rationality." In *Encyclopedia of Law and Economics*, edited by Alain Marciano and Giovanni Battista Ramello. Springer New York. https://doi.org/10.1007/978-1-4614-7753-2_404.

Newell, Allen, and Herbert A Simon. 1976a. "Computer Science as Empirical Inquiry: Symbols and Search." *Communications of the ACM* 19 (3): 113–26.

Newell, Allen, and Herbert A. Simon. 1976b. "Computer Science as Empirical Inquiry: Symbols and Search." *Commun. ACM* (New York, NY, USA) 19 (3): 113–26.

<https://doi.org/10.1145/360018.360022>.

Nilsson, Nils J. 2005. "Human-Level Artificial Intelligence? Be Serious!" *AI Mag.* 26 (4): 68–75. <https://doi.org/10.1609/AIMAG.V26I4.1850>.

Pearl, Judea. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.

Russell, Stuart, and Peter Norvig. 2020. *Artificial Intelligence: A Modern Approach*. 4th ed. Pearson. <http://aima.cs.berkeley.edu/>.

Savage, Neil. 2024. "Beyond Turing: Testing LLMs for Intelligence." *Commun. ACM* (New York, NY, USA), ahead of print, June. <https://doi.org/10.1145/3673427>.

Strickland, Eliza. 2021. "The Turbulent Past and Uncertain Future of AI: Is There a Way Out of AI's Boom-and-Bust Cycle?" *IEEE Spectrum* 58 (10): 26–31. <https://doi.org/10.1109/MSPEC.2021.9563956>.

Wang, Pei. 2019. "On Defining Artificial Intelligence." *Journal of Artificial General Intelligence* 10 (2): 1–37. <https://doi.org/10.2478/jagi-2019-0002>.

Appendix: On Defining Artificial Intelligence

Wang (2019)

An **agent** and its **interaction with the environment** are specified as a tuple:

$$\langle P, S, A \rangle$$

where - P represents a sequence of **input signals**, $P = \langle p_0, \dots, p_t \rangle$ - S represents a sequence of **internal states**, $S = \langle s_0, \dots, s_t \rangle$ - A represents a sequence of **actions**, $A = \langle a_0, \dots, a_t \rangle$ For a sequence of moments, $0, \dots, t$. ## Human (H) vs Computer (C) > **(Wang 2019)** > > AI is conceived as computer systems that are similar to the human mind > in a certain sense, though a computer and a human mind cannot be > identical in all aspects.

$$\langle P^H, S^H, A^H \rangle \approx \langle P^C, S^C, A^C \rangle$$

. . . Wang (2019) proposes 5 perspectives: **Structure-AI**, **Behavior-AI**, **Capability-AI**, **Function-AI**, and **Principle-AI**. ## 1. Structure-AI (brain modelling, cognitive science) > **(Wang 2019)** > > I call this type of definition "**Structure-AI**," since it requires > an AI system to go through **isomorphic states** or structure changes > as the brain does when

they are given **similar input**, which will > produce **similar output**, so the three components of the two are > pairwise similar to each other:

$$P^H \approx P^C, S^H \approx S^C, A^H \approx A^C$$

2. Behaviour-AI (Turing Test) > **(Wang 2019)** > > One way to acknowledge a human-like mind without demanding a > human-like brain is to associate intelligence to the external > behaviors of the agent. After all, **if an agent behaves like a human, > it should be considered as intelligent**, no matter whether it looks > like a human, either inside or outside.

$$P^H \approx P^C, A^H \approx A^C$$

3. Capability-AI (Employment Test) > **(Wang 2019)** > > In the agent framework, it means that C is similar to H in the > sense that there are moments i and j that:

$$p_i^C \approx p_j^H, a_i^C \approx a_j^H$$

> **(Wang 2019)** > > the action (solution) the computer produces for a percept (problem) is > similar to the action produced by a human to a similar percept > (. . .) In this way, **the intelligence of a system is identified > by a set of problems it can solve, while whether they are solved > in the "human way" does not matter.** ## Capability-AI (contd) > **(Nilsson 2005)** > > "I suggest we replace the Turing test by something I will call the > 'employment test.' To pass the employment test, **AI programs must be > able to perform the jobs ordinarily performed by humans.** Progress > toward human-level AI could then be measured by the fraction of these > jobs that can be acceptably performed by machines" ## 4. Function-AI > **(Wang 2019)** > > In the agent framework, this **"Function-AI"** perspective takes C > to be similar to H in the sense that there are moments i and j > that:

$$a_i^C \approx f^C(p_i^C), a_j^H \approx f^H(p_j^H), f^C \approx f^H$$

> **(Wang 2019)** > > Here the function can correspond to searching, **reasoning**, > **learning**, etc., and since **the focus is on the functions** (i.e., > input-output mappings), the concrete input and output values of the > two agents do not have to be similar to each other. ## 5. Principle-AI (rationality, logicist) > **(Wang 2019)** > > As in any field, there are researchers in AI trying to find > **fundamental principles** that can uniformly explain the relevant > phenomena. Here the idea comes from the usage of intelligence as a > form of *rationality* (. . .) that can make the **best-possible > decision** in various situations, according to the experience or > history of the system.

$$A^C = F^C(P^C), A^H = F^H(P^H), F^C \approx F^H$$

(Wang 2019)

The above F is often not formally specified, but described informally as a certain “**principle,**” which is not merely about a single type of problem and its solution, but about **the agent’s life-long history in various situations,** when dealing with various types of problems.

Code of the day

```
#!/usr/bin/env python3
# -*- Mode: Python -*-
# ai_lecture-01.py
# Author          : Marcel Turcotte & ChatGPT 5
# Created On      : Tue Feb 13 16:29:41 2024
# Last Modified By: Marcel Turcotte
# Last Modified On: Sun Sep  6 09:58:08 EDT 2026

# The initial version of this script was developed in 2024.
# In 2025, ChatGPT was used to revise the code to align with the
# latest API version, expand the comments to improve its
# instructional value,
# and enhance its suitability as a teaching example. In 2026, the
# script was
# further updated using Codex ChatGPT 5.6 Sol.
# Attempts to use Codex ChatGPT 6 Astra Light were disappointing.
```

```
"""Generate bilingual narration and speech for the first lecture.
```

```
The script uses the OpenAI API to write an English narration,
translate it
into Canadian French, and synthesize both versions as audio files.
```

```
Required:
```

```
python -m pip install openai
```

```
OPENAI_API_KEY can be exported by the shell or injected by a secret
manager:
```

```
export OPENAI_API_KEY="your_api_key_here"
python ai_lecture-01-v04.py --voice marin
```

```
With a 1Password environment file containing a secret reference:
```

```
op run --env-file="$HOME/.env" -- python ai_lecture-01-v04.py -
--voice marin
```

```
Optional .env support:
```

```
python -m pip install python-dotenv
```

```
Keep API keys private and exclude .env files from version control.
When presenting the recordings, disclose that the voices are AI-
generated.
```

```
"""
```

```

from __future__ import annotations

import argparse
import os
import sys
from pathlib import Path
from typing import Optional

from openai import (
    APIConnectionError,
    APITimeoutError,
    AuthenticationError,
    OpenAI,
    OpenAIError,
    RateLimitError,
)

try:
    from dotenv import load_dotenv
except ImportError:
    load_dotenv = None

TEXT_MODEL = "gpt-4o"
SPEECH_MODEL = "gpt-4o-mini-tts"

WELCOME_SENTENCE = 'Welcome to CSI 4106, "Introduction to  
Artificial Intelligence"!'

# Quoted verbatim from the official uOttawa calendar.
OFFICIAL_CALENDAR_DESCRIPTION = (
    "The roots and scope of Artificial Intelligence. Knowledge and  
knowledge "  
    "representation. Search, informed search, adversarial search.  
Deduction and "  
    "reasoning. Uncertainty in Artificial Intelligence.  
Introduction to Natural "  
    "Language Processing. Elements of planning. Basics of Machine  
Learning."  

)

TEACHING_APPROACH = (
    "The calendar description establishes the broad scope of the  
course rather than its "  
    "teaching order. This offering uses a modernized, machine-  
learning-first structure and "  
    "introduces deep learning as early as the necessary foundations  
allow. Deep learning "  
    "is central to contemporary AI, so studying it early connects  
the course with "  
    "technologies that attract students' attention. It also  
provides a framework for "
)

```

```

        "introducing and defining central AI ideas. We will examine
        cases in which deep "
        "learning has displaced or changed the role of earlier
        approaches, as well as cases "
        "in which its limitations make earlier, purpose-built methods
        valuable. A further "
        "motivation is to treat learning as an early and extensively
        studied milestone in the "
        "development of intelligence. This foundation will later help
        us compare learning-based "
        "and search-based approaches when studying Monte Carlo Tree
        Search."
    )

```

```

LEARNING_OUTCOMES = (
    "Explain the fundamental concepts of Artificial Intelligence
    (AI)",
    "Apply problem-solving strategies using AI techniques",
    "Critically analyze and compare different AI approaches",
    "Demonstrate independent learning and exploration",
)

```

```

PRELIMINARY_COURSE_OUTLINE = (
    "Defining AI",
    "Introduction to machine learning",
    "Linear regression",
    "Training",
    "Logistic regression",
    "Evaluation",
    "Hyperparameter tuning",
    "Machine learning engineering",
    "Introduction to artificial neural networks",
    "Backpropagation",
    "Softmax, cross-entropy, and regularization",
    "Convolutional neural networks",
    "Introduction to search",
    "Informed search",
    "Local search",
    "Population-based metaheuristics",
    "Adversarial search",
    "Monte Carlo Tree Search",
)

```

```

# -----
# Setup
# -----

```

```

def prepare_api_key() -> bool:
    """Load an optional .env file and check for an API key."""
    if not os.getenv("OPENAI_API_KEY") and load_dotenv is not None:
        load_dotenv(override=False)

```

```

api_key = os.getenv("OPENAI_API_KEY", "")

if api_key.startswith("op://"):
    print(
        "OPENAI_API_KEY is an unresolved 1Password reference.
Run:\n"
        f'  op run --env-file="$HOME/.env" -- python
{Path(__file__).name}',
        file=sys.stderr,
    )
    return False

if not api_key:
    print(
        "No OpenAI API key was found. You can:\n"
        "  • export OPENAI_API_KEY=\"your_api_key_here\"\n"
        "  • use a secret manager to inject OPENAI_API_KEY\n"
        "  • install python-dotenv and store the key in a local
.env file\n\n"
        "An exported key or a secret manager does not require
python-dotenv.",
        file=sys.stderr,
    )
    return False

return True

# -----
# Prompt builders
# -----

def build_course_context() -> str:
    """Organize the source material while keeping its roles
unambiguous."""
    outcomes = "\n".join(f"- {outcome}" for outcome in
LEARNING_OUTCOMES)
    outline = "\n".join(
        f"{number}. {topic}"
        for number, topic in enumerate(PRELIMINARY_COURSE_OUTLINE,
start=1)
    )

    return f'''REQUIRED OPENING SENTENCE:
{WELCOME_SENTENCE}

OFFICIAL UOTTAWA CALENDAR DESCRIPTION:
{OFFICIAL_CALENDAR_DESCRIPTION}

INSTRUCTOR'S DESIGN FOR THIS OFFERING:
{TEACHING_APPROACH}

```

LEARNING OUTCOMES:

Upon completion of the course, students will be able to:
{outcomes}

PRELIMINARY AND AMBITIOUS COURSE OUTLINE:

{outline}'''

```
def build_instructions_en(tone: Optional[str]) -> str:
    """Build instructions for the English narration."""
    instructions = (
        "Write a natural spoken welcome of 180 to 240 words for the  
first lecture of an "  
        "undergraduate AI course at the bilingual University of  
Ottawa. Begin with the "  
        "required opening sentence exactly as supplied. Clearly  
identify the calendar "  
        "paragraph as the official uOttawa course description, then  
explain that this "  
        "offering uses a modernized, machine-learning-first  
structure and reaches deep "  
        "learning as early as its foundations permit. Do not  
present the calendar "  
        "description as the teaching sequence or the preliminary  
outline as official "  
        "calendar text. Explain both reasons for teaching deep  
learning early: its "  
        "importance to current AI and its use as a framework for  
discussing other AI "  
        "ideas. Contrast cases where deep learning has changed  
earlier approaches with "  
        "cases where its limitations make specialized methods  
valuable. State all four "  
        "learning outcomes naturally, preserving their action  
verbs. Describe the "  
        "preliminary and ambitious sequence as an arc from machine-  
learning foundations "  
        "through neural networks and then search, ending with Monte  
Carlo Tree Search; do "  
        "not recite all eighteen topics. Do not imply that deep  
learning is a prerequisite "  
        "for Monte Carlo Tree Search. The students have varied  
backgrounds, languages, and "  
        "prior exposure to AI. Be welcoming, inclusive, clear, and  
encouraging. Use plain "  
        "language and short sentences. Return only the narration."  
    )
    if tone:
        instructions += f" Writing style: {tone}"
    return instructions
```

```
def build_instructions_fr() -> str:
    """Build instructions for the Canadian French translation."""
    return (
        "Translate the supplied English narration into natural  

        Canadian French. "  

        "The University of Ottawa is bilingual, and CSI 4506 is the  

        French-language "  

        "course corresponding to CSI 4106. Replace 'CSI4106' or  

        'CSI 4106' with "  

        "'CSI 4506'. Translate the course title as 'Introduction à  

        l'intelligence "  

        "artificielle'. Preserve the distinction between the  

        official calendar description "  

        "and the way this offering is organized, along with all  

        other ideas and the "  

        "welcoming tone. Return only the translation."
    )
```

```
def build_speech_instructions(language: str, tone: Optional[str]) -> str:
    """Build language-specific delivery instructions."""
    if language == "en":
        instructions = (
            "Speak as a welcoming university instructor. Use  

            natural Canadian English, "  

            "a conversational tone, and a measured pace. Pause  

            briefly between the main "  

            "ideas. Pronounce 'CSI 4106' as 'C S I four one zero  

            six'."
        )
    else:
        instructions = (
            "Parlez comme une personne qui accueille sa classe à  

            l'université. Utilisez "  

            "un français canadien naturel, un ton chaleureux et un  

            débit mesuré. Faites "  

            "une courte pause entre les idées principales.  

            Prononcez « CSI 4506 » comme "  

            "« cé esse i, quatre cinq zéro six »."
        )

    if tone:
        instructions += f" Additional delivery instruction: {tone}"
    return instructions
```

```
# -----
# OpenAI utilities
# -----
```

```

def generate_text(client: OpenAI, source: str, instructions: str) -
> str:
    """Generate text with the Responses API."""
    response = client.responses.create(
        model=TEXT_MODEL,
        instructions=instructions,
        input=source,
        temperature=0.2,
        max_output_tokens=800,
    )
    return (response.output_text or "").strip()

```

```

def synthesize_speech(
    client: OpenAI,
    text: str,
    output_path: Path,
    *,
    voice: str,
    instructions: str,
    response_format: str,
) -> None:
    """Stream synthesized speech to an audio file."""
    with client.audio.speech.with_streaming_response.create(
        model=SPEECH_MODEL,
        voice=voice,
        input=text,
        instructions=instructions,
        response_format=response_format,
    ) as response:
        response.stream_to_file(output_path)

```

```

# -----
# Script logic
# -----

```

```

def main(audio_format: str, voice: str, tone: Optional[str]) ->
int:
    """Generate English and French narrations and audio files."""
    if not prepare_api_key():
        return 1

    course_context = build_course_context()

    try:
        client = OpenAI()

        narration_en = generate_text(
            client, course_context, build_instructions_en(tone)
        )
        if not narration_en:

```

```

        print("The text model returned an empty English
narration.", file=sys.stderr)
        return 1

    narration_fr = generate_text(
        client, narration_en, build_instructions_fr()
    )
    if not narration_fr:
        print("The text model returned an empty French
translation.", file=sys.stderr)
        return 1

    narrations = (
        ("en", narration_en),
        ("fr", narration_fr),
    )

    for language, narration in narrations:
        output_path = Path(
            f"01_tts_course_intro-{language}-{voice}.
{audio_format}"
        )
        synthesize_speech(
            client,
            narration,
            output_path,
            voice=voice,
            instructions=build_speech_instructions(language,
tone),
            response_format=audio_format,
        )
        print(f"[OK] {language.upper()} audio → {output_path}")

    except AuthenticationError:
        print(
            "OpenAI rejected the API key. Check OPENAI_API_KEY or
the secret manager.",
            file=sys.stderr,
        )
        return 1
    except (APIConnectionError, APITimeoutError) as error:
        print(f"Could not reach the OpenAI API: {error}",
file=sys.stderr)
        return 1
    except RateLimitError:
        print(
            "The request was rate-limited. Check the account's
limits and API credit.",
            file=sys.stderr,
        )
        return 1
    except OpenAIError as error:

```

```

        print(f"OpenAI API error: {error}", file=sys.stderr)
        return 1
    except OSError as error:
        print(f"Could not write an audio file: {error}",
file=sys.stderr)
        return 1

    print("Reminder: disclose that the narration uses AI-generated
voices.")
    return 0

# -----
# Command-line interface
# -----

if __name__ == "__main__":
    parser = argparse.ArgumentParser(
        description="Generate bilingual first-lecture narration and
speech."
    )
    parser.add_argument(
        "--format",
        default="mp3",
        choices=("mp3", "wav", "aac", "flac", "opus", "pcm"),
        help="Audio output format (default: mp3).",
    )
    parser.add_argument(
        "--voice",
        default="marin",
        help="Speech voice (default: marin). Try cedar or nova.",
    )
    parser.add_argument(
        "--tone",
        default=None,
        help=(
            "Optional writing and delivery instruction, for
example: "
            "'Sound warm, calm, and optimistic.'"
        ),
    )
    args = parser.parse_args()

    raise SystemExit(main(args.format, args.voice, args.tone))

```

Marcel **Turcotte**

Marcel.Turcotte@uOttawa.ca

School of Electrical Engineering and **Computer Science (EECS)**

