Lecture 01.05 Robot intelligence

Intelligence is famously challenging to define, but that won't stop us.

Definition 01.05.1: intelligence

Intelligence is the ability to map perceptions to performant actions.

Let's consider this definition specifically with regard to *robot intelligence*, a form of *artificial intelligence*. Here, *perceptions* are measurements, discussed in Lecture 01.03. *Actions* are with the robot's environment (Lecture 01.04). The remaining qualifier here is *performant*: to perform well with respect to a *metric*.

robot intelligence
artificial intelligence
perceptions
actions
performant
metric

Let's consider this last requirement in more detail, with reference to Figure 01.7. If a robot action map did not have a metric, it would offend our sense of the meaning of the term "intelligence." For instance, if a vehicle in locomotion perceived an obstacle, was mechanically capable of avoiding a collision, but did not, it would not seem intelligent. If one of its metrics is collision-avoidance, it performed poorly and looks dumb. If it didn't have such a metric in the first place, it looks not-even-dumb. Stupidly, it was lacking an obvious aspect of locomotion.

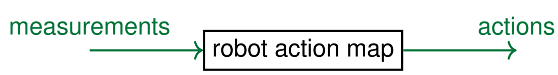

Figure 01.7: robot action map.

The term *goal* or *objective* is frequently used here, in the sense that for something to be intelligent it must have one. But there is something more implied in the term "goal": *intention*. It was the study of biology that first alerted us to the limitations of requiring intelligence to involve intention. Many biological systems exhibit what we would call intelligence, such as the flocking of birds as they fly together, without any sort of coordinated effort. In fact, flocking behavior can be replicated in groups of robots based on simple rules given to each robot. That is, a complex behavior that could be assigned a metric (e.g. flocking could be measured by the efficiency of group flying) can emerge without intention. This is called *emergence*, a topic that is perhaps more mysterious than is necessary.

goal
objective
intention
emergence

01.05.1 Robotic applications of artificial intelligence

Before we introduce some of the methods of artificial intelligence, it is worth considering some of its applications in robotics.

Machine vision systems process image sensor data streams to construct actionable information. This includes object detection and internal model construction.

SLAM or simultaneous localization (of the robot) and mapping (of its environment) must fuse (sensor fusion!) vision and other sensor data. AI can contribute to this process.

Path planning is the process of planning the route the robot should take through an environment. For instance, self-driving cars use map data and GPS to plan its route.

Natural language processing can make the robot responsive to human speech and allow it to construct human-understandable speech.

01.05.2 Methods of robot intelligence

The development of artificial intelligence (AI) has been so varied in method that summaries invariably and unjustly minimize entire fields of research. Therefore, instead of attempting a hierarchical organization, we will highlight a few ideas of particular interest.

symbolic manipulation

The first is that of *symbolic manipulation*. Here intelligence is understood as the use of symbolic representations and rules of inference. For instance,

> *if there is a parking spot available on the street, park;*
> *otherwise, circle the block.*

signal-to-symbol problem

natural language processing

There are several symbols here, e.g. a parking spot, the street, parking, block, circling. This approach has encountered formidable challenges, such as the *signal-to-symbol problem* of going from measurement signals to appropriate symbolic representations. In the example, it is challenging to determine from sensor measurements what is a parking spot. This approach also encounters the vast ambiguities in natural language; consider the multiple potential meanings of the sentence, "There's a spot." *Natural language processing* is now an entire field of AI.

logical reasoning
knowledge base
semantics
knowledge representations
planning
probabilistic reasoning

Logical approaches to AI understand intelligence to be comprised of *reasoning*. Reasoning requires a *knowledge base* with *semantic* (i.e. meaningful) relations, so *knowledge representations* have been a significant aspect of this work. This approach features, for instance, the ability to make *plans*. In reality, nothing is certain, so logical methods have come to use *probabilistic*

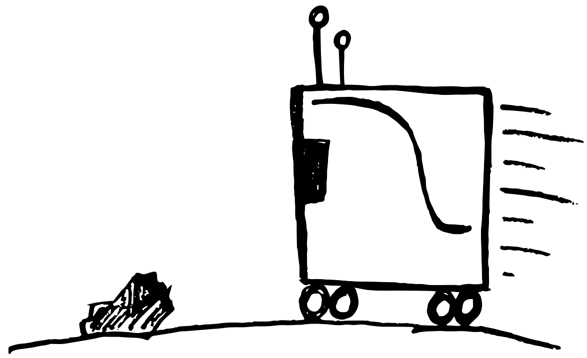

Figure 01.8: a robotic vehicle hurtling toward an obstacle. (PR)

reasoning with, for instance, *Bayesian networks*. Another approach to uncertain reasoning is *fuzzy logic* in which elements can be in a set to a degree. **Bayesian networks** **fuzzy logic**

Beyond the minimum requirement set forth in our definition of intelligence, it is certainly a mark of intelligence to *adapt* or *learning*: to increase intelligence from experience. So, consider again the vehicle in locomotion and perceiving an obstacle it is mechanically capable of avoiding, now illustrated in Figure 01.8. If it has as a metric to avoid obstacles, it might still hit the obstacle if its action map is insufficiently intelligent. Now consider if this map is *trained* on many obstacles such that it tweaks the map in accordance with its obstacle-avoidance metric. This would yield an action map the performance of which improves in intelligence. Taken broadly, this program is called *machine learning*, a leading strategy for developing artificial intelligence. **adaptation** **learn** **training** **machine learning**

One method of machine learning uses large amounts of data to make inferences about what should be done, making it inherently *statistical*. A simple example here is the technique of *regression*: fitting a function to data by optimizing the function's parameters. **statistics** **regression**

Another machine learning approach is to train an *artificial neural network*, a mathematical operation that approximates the function of biological neural networks that are key to brain activity. Typically, a set of *training data* that includes *labels* (the "correct" answers, usually assigned by humans) is processed by the network, which adapts its parameters such that its output approaches the "correct" answers. Then labeled *testing data*, on which the network has *not* trained, tests to see how well the network performs on new data. **neural network** **training data** **labels** **testing data**

Relatively recently, a combination of advances in computing power and new approaches to using neural networks has been highly successful, surpassing previous performance metrics for many important applications (e.g. vision, natural language processing, game playing, etc.). The new techniques use multiple network *layers* and are collectively called *deep learning* (deep in layers).

layers
deep learning

01.05.3 Control theory

control system
control theory
plant
feedback
controller

A *control system*, which *control theory* studies, is comprised of a system to be controlled called the *plant*, a *feedback* measurement of the plant's output, and a *controller* that determines the plant's input given the feedback and some goal output. This is frequently represented in the block diagram of Figure 01.9.

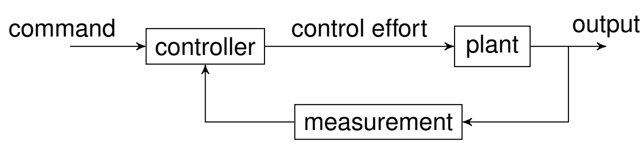

Figure 01.9: a feedback control system block diagram.

If it is successful at meeting performance goals, such a system meets all the requirements of artificial intelligence in our definition. Although control theory developed as a separate field from AI, it is in fact another form and is now generally recognized as such.

Control systems appear all over most robots as subsystems. For instance, a robot arm joint usually includes a motor (actuator), an encoder (sensor), and a controller. The plant, then, might be the motor-link assembly. A process in some higher-level controller that perhaps has planned the arm's motion and performed an inverse kinematics calculation would send a command to the joint to rotate to a specific angle.

01.05.4 General artificial intelligence

general AI
strong AI
weak AI

There may be objections at this point about our threshold for intelligence being too low. Some find it to so because they conceive of intelligence as being what is now called *general artificial intelligence*: a single AI system applicable to *any* problem. For instance, it would be capable of navigating a robot, having a conversation, and playing the cello. This is sometimes called *strong AI*, in contrast with the version considered above, now termed *weak AI*.

There is no strong AI, at present. However, if it does emerge, it is possible that it will develop beyond human intelligence rather quickly, into a state called *superintelligence*. There are ethical concerns here, with some potential for superintelligence becoming an existential threat to the human species.

superintelligence