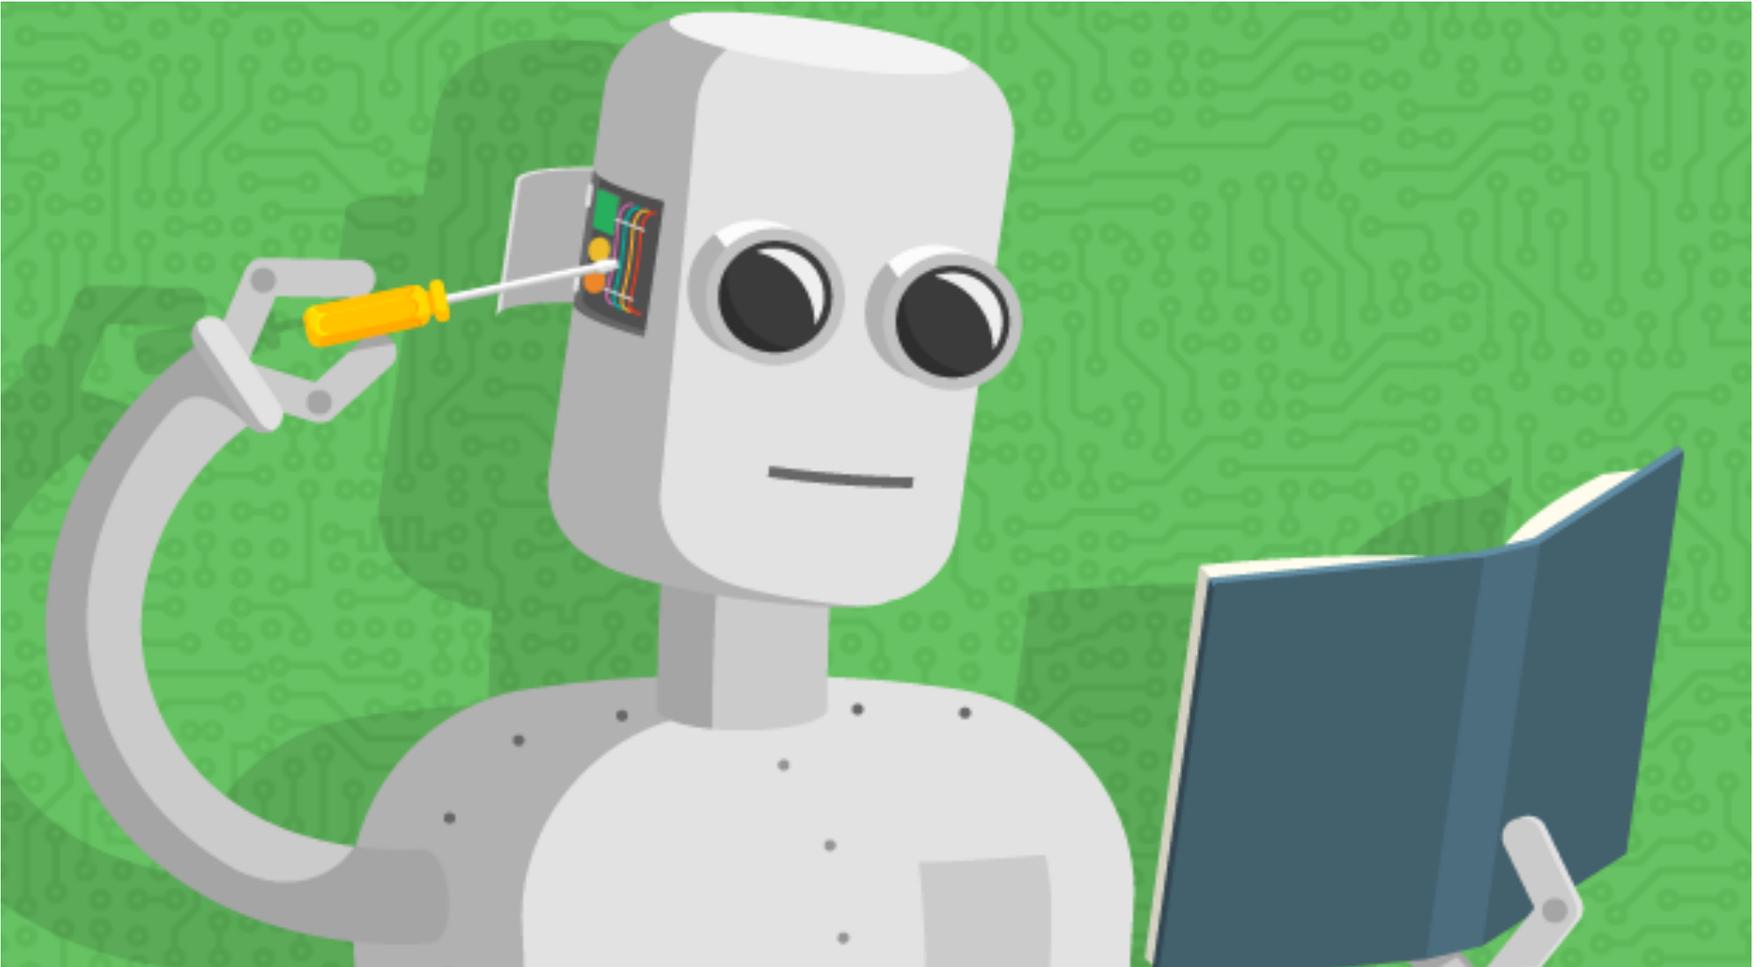


Machine Learning overview

Chapter 18, 21



What is learning?

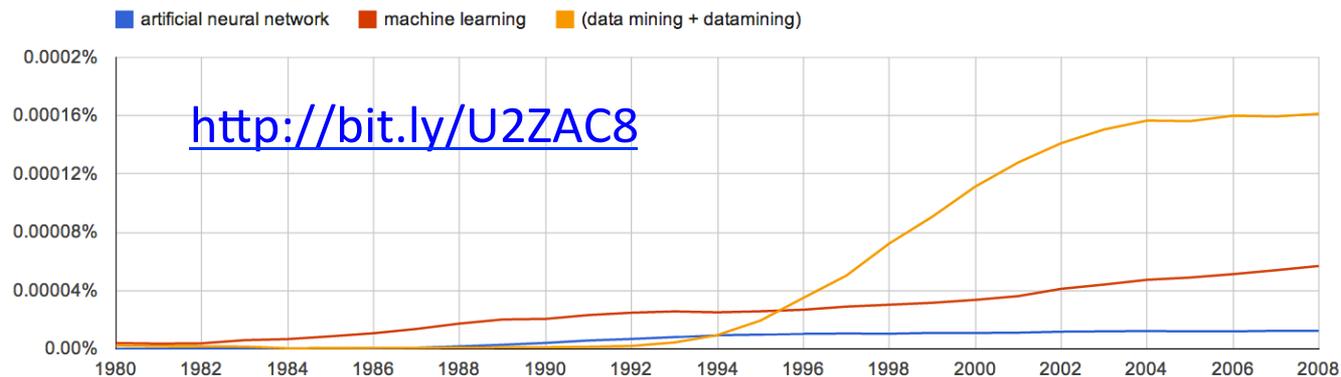
- “Learning denotes changes in a system that ... enable a system to do the same task more efficiently the next time” – [Herbert Simon](#)
- “Learning is constructing or modifying representations of what is being experienced” – [Ryszard Michalski](#)
- “Learning is making useful changes in our minds” – [Marvin Minsky](#)

Why study learning?

- Understand and improve efficiency of **human learning**
 - Use to improve methods for teaching and tutoring people (e.g., better computer-aided instruction)
- **Discover** new things or structure previously unknown
 - Examples: data mining, scientific discovery
- Fill in skeletal or **incomplete specifications in** a domain
 - Large, complex systems can't be completely built by hand & require dynamic updating to incorporate new information
 - Learning new characteristics expands the domain or expertise and lessens the “brittleness” of the system
- Build agents that can **adapt** to users, other agents, and their environment

AI & Learning Today

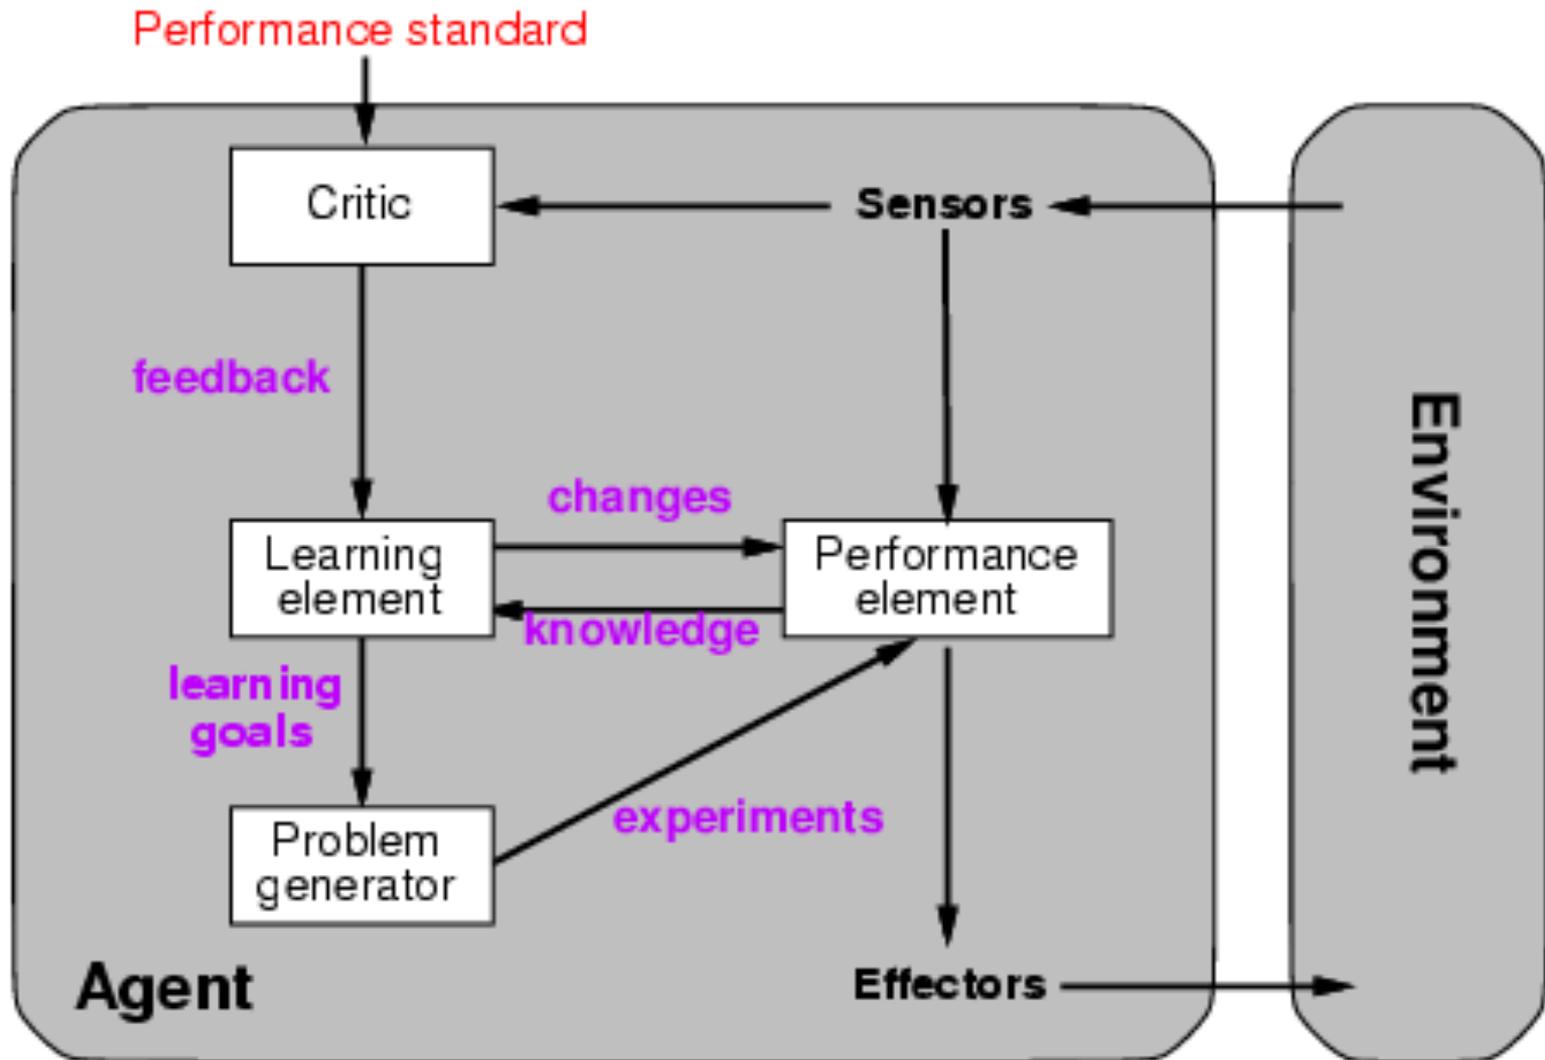
- Neural network learning was popular in the 60s
- In the 70s and 80s it was replaced with a paradigm based on manually encoding and using knowledge
- In the 90s, more data and the Web drove interest in new statistical machine learning (ML) techniques and new data mining applications
- Today, ML techniques and big data are behind almost all successful intelligent systems



Machine Learning Successes

- Sentiment analysis
- Spam detection
- Machine translation
- Spoken language understanding
- Named entity detection
- Self driving cars
- Motion recognition (Microsoft X-Box)
- Identifying faces in digital images
- Recommender systems (Netflix, Amazon)
- Credit card fraud detection

A general model of learning agents



Major paradigms of machine learning

- **Rote learning:** 1-1 mapping from inputs to stored representation, learning by memorization, association-based storage and retrieval
- **Induction:** Use specific examples to reach general conclusions
- **Clustering:** Unsupervised discovery of natural groups in data
- **Analogy:** Find correspondence between two different representations
- **Discovery:** Unsupervised, specific goal not given
- **Genetic algorithms:** *Evolutionary* search techniques, based on an analogy to *survival of the fittest*
- **Reinforcement** – Feedback (positive or negative reward) given at the end of a sequence of steps

What we will and won't cover

- We'll look at a few popular machine learning problems and algorithms
 - Take CMSC 478/678 Machine Learning for more
 - Use online resources & experiment on your own
- We'll focus on when/how to use techniques and only touch on how/why they work
- We'll cover basic methodology and evaluation
- We'll use the [Weka](#) platform for examples and demonstrations
 - Great for exploration and learning

Types of learning problems

- **Supervised:** learn from training examples
 - Regression
 - Classification: Decision Trees, SVM
- **Unsupervised:** learn w/o training examples
 - Clustering
 - Word semantics
- Lots more we won't cover
 - Hidden Markov models
 - Learning to rank
 - Semi-supervised learning
 - Reinforcement learning
 - Active learning

Supervised learning

- Given training examples of inputs and corresponding outputs, produce the “correct” outputs for new inputs
- Two main scenarios:
 - **Classification:** outputs typically labels (goodRisk, badRisk); learn a decision boundary that separates classes
 - **Regression:** also known as “curve fitting” or “function approximation.” Learn a continuous input-output mapping from examples (possibly noisy)

Unsupervised Learning

- Given only *unlabeled* data as input, learn some sort of structure
- E.g., cluster your Facebook friends based on similarity of posts
- E.g., find sets of words whose meanings are related

Weka

Open source Java software for ML and datamining

<http://cs.waikato.ac.nz/ml/weka/>

The screenshot displays the Weka Explorer interface. At the top, there are tabs for 'Preprocess', 'Classify', 'Cluster', 'Associate', 'Select attributes', and 'Visualize'. Below these are buttons for 'Open file...', 'Open URL...', 'Open DB...', 'Generate...', 'Undo', 'Edit...', and 'Save...'. The 'Filter' section shows a 'Choose' button and a text field containing 'Discretize -B 3 -M -1.0 -R 1'. The 'Current relation' section shows 'auto-mpg' with 'Attributes: 8' and 'Sum of weights: 240'. A table of statistics for the 'mpg' attribute is shown:

Statistic	Value
Minimum	9
Maximum	44.6
Mean	23.006
StdDev	7.777

Below the statistics is a 'Class: origin (Nom)' dropdown and a 'Visualize All' button. A histogram shows the distribution of 'mpg' values, with bars colored in blue and red. The x-axis ranges from 9 to 44.6, and the y-axis shows counts for each bin: 24, 50, 49, 39, 35, 25, 14, 4.

A yellow tooltip points to the 'mpg' attribute in the list, stating: 'Left-click to edit properties for this object, right-click/Alt+Shift+left-click for menu'.

The screenshot shows the 'Weka GUI Chooser' dialog box. It features a window title 'Weka GUI Chooser' and a menu bar with 'Program', 'Visualization', 'Tools', and 'Help'. On the left, there is a logo of a weka bird and the text 'WEKA The University of Waikato' and 'Weka, a native bird of New Zealand'. Below this is the version information: 'Waikato Environment for Knowledge Analysis Version 3.8.0 (c) 1999 - 2016 The University of Waikato Hamilton, New Zealand'. On the right, under the 'Applications' section, there are five buttons: 'Explorer', 'Experimenter', 'KnowledgeFlow', 'Workbench', and 'Simple CLI'.