













### What is learning?

Definition: "... changes to the content and organization of an agent's knowledge enabling it to improve its performance on a particular task or population of tasks" [Herb Simon].











## Unsupervised

- unsupervised learning need not involve a teacher at all.
- if it does teacher only provides rewards and punishments.
- The learner not only has to figure out <u>what</u> constitutes a problem instance but also has to figure out an appropriate <u>response</u>
- e.g., learn chess by trial and error or discovery learning as in learning new concepts in mathematics (perfect numbers) or game playing (in simulated sea battles lots of small cruisers are often better than a few very large battleships).







# Example presentation

- Batch problems all of the examples are given at once
- Online problems, the learner is given one example at a time and is assumed not to have the storage necessary to keep track of all of the examples seen so far.
- Why is online interesting? Examples?











## **Inductive Bias: definition**

- This "some idea of what to choose from" is called an **inductive bias**.
- Terminology
  - H, **hypothesis space** a set of functions to choose from C, **concept space** a set of possible functions to learn
- Often in learning we search for a hypothesis f in H that is consistent with the training examples, i.e.,
  f(x) = y for all training examples (x,y).
- In some cases, any hypothesis consistent with the training examples is likely to generalize to unseen examples. The trick is to find the right bias.