Learning Agents

Overview

Learning

important aspects

Learning in Agents

goal, types; individual agents, multi-agent systems

Learning Agent Model

components, representation, feedback, prior knowledge

Learning Methods

inductive learning, neural networks, reinforcement learning, genetic algorithms

Knowledge and Learning

explanation-based learning, relevance information

Learning

acquisition of new knowledge and skills
on the agent's own initiative

incorporation of new knowledge
into the existing knowledge

performed by the system itself
not only injected by the developer

performance improvement
simply accumulating knowledge isn't sufficient

Learning in Agents

improved performance through learning

learning

modify the internal knowledge

goal

improvement of future performance

types of learning

memorization, self-observation, generalization,
exploration, creation of new theories,
meta-learning

levels of learning

value-action pairs
representation of a function
general first-order logic theories

Learning Agent Model

conceptual components

learning element

responsible for making improvements

performance element

selection of external actions:

takes in percepts and decides on actions

critic

evaluation of the performance according to a fixed standard

problem generator

suggests exploratory actions

new experiences with potential benefits

Diagram [**?**] p. 526

Learning Element

how to improve performance

performance element

affected components

internal representation

used for components to improved

feedback

from the environment

from a teacher

prior knowledge

about the environment / domain

Performance Element Components

relevant for learning

mapping function
: from percepts and internal state to actions

inference mechanism
: infer relevant properties of the world from percepts

changes in the world
: information about the way the world evolves

effects of actions
: results of possible actions the agent can take

utility information
: desirability of world / internal states

action-value information
: desirability of actions in particular states

goals
: classes of desirable states

utility maximization

Representation

used in a component

deterministic

linear weighted polynomials

logic

propositional, first order

probabilistic

belief networks, decision theory

learning algorithms need to be adapted to the particular representation

Feedback

about the desired outcome

supervised learning

inputs and outputs of percepts can be perceived

immediately

reinforcement learning

an evaluation of the action (hint) becomes

available

not necessarily immediately

no direct information about the correct action

unsupervised learning

no hint about correct outputs

Inductive Learning

learning from examples

reflex agent

direct mapping from percepts to actions

inductive inference

given a collection of examples for a function f,
return a function h (hypothesis) that
approximates f

bias

preference for one hypothesis over another
usually large number of possible consistent
hypotheses

incremental learning

new examples are integrated as they arrive

Decision Trees

deriving decisions from examples

goal

take a situation described by a set of properties,] and produce a yes/no decision

goal predicate

Boolean function defining the goal

expressiveness

propositional logic

efficiency

more compact than truth tables in many cases

exponential in some cases (parity, majority)

Induction

for decision trees

example

described by the values of the attributes and
the value of the goal predicate (classification)

training set

set of examples used for training

test set

set of examples used for evaluation
different from the training set

algorithm

classify into positive and negative sets
select the most important attribute
split the tree, and apply the algorithm
recursively to the subtrees

Performance Evaluation

for inductive learning algorithms

goals

reproduce classification of the training set

predict classification of unseen examples

example set size

must be reasonably large

average prediction quality

for different sizes of training sets and randomly selected training sets

learning curve ("happy curve")

plots average prediction quality as a function of the size of the training set

training and test data should be kept separate, and each run of the algorithm should be independent of the others

Examples

decision tree learning

GASOIL

design of oil platform equipment

expert system with 2500 rules generated from existing designs

using a flight simulator

program generated from examples of skilled human pilots

somewhat better performance that the teachers (for regular tasks)

not so good for rare, complex tasks

Neural Networks: see separate slides

Reinforcement Learning

learning from success and failure

reinforcement or punishment

- feedback about the outcome of actions
- no direct feedback about the correctness of an action
- possibly delayed

rewards as percepts

- must be recognized as special percepts, not just another sensory input
- can be components of the utility, or hints

Variations

in the learning task

environment

accessible or not

prior knowledge

internal model of the environment

knowledge about effects of actions

utility information

passive learner

watches the environment without actions

active learner

act based upon learned information

problem generation for exploring the environment

exploration

trade-off between immediate and future benefits

Generalization

in reinforcement learning

implicit representation

more compact form than a table for input-output values

input generalization

apply learned information to unknown states

trade-off between the size of the hypothesis space and the time to learn a function

Examples

of reinforcement learning

game-playing

TD-gammon: neural network with 80 hidden units, 300,000 training games and precomputed features added to the input representation plays on par with the top three human players worldwide

robot control

cart-pole balancing (inverted pendulum)

Genetic Algorithms

as a variation of reinforcement learning

basic idea

selection and reproduction operators are applied to sets of individuals

reward

successful reproduction

agent is a species, not an individual

fitness function

takes an individual, returns a real number

algorithm

parallel search in the space of individuals for one that maximizes the fitness function

selection strategy

random, probability of selection is proportional to fitness

reproduction

selected individuals are randomly paired

cross-over: gene sequences are split at the same point and crossed

mutation: each gene can be altered with small probability

Knowledge and Learning

learning with prior knowledge

learning methods

take advantage of prior knowledge about the environment

learning level

general first-order logic theories
as opposed to function learning

description

conjunction of all example specifications

classification

conjunction of all example evaluations

hypothesis

newly generated theory

entailment constraint

together with descriptions, the hypothesis must entail classifications

essential step toward truly autonomous intelligent agents

Explanation-Based Learning

potentially known information is made explicit

usage

known theories are converted into directly applicable knowledge ("aha-effect")

entailment constraint

background entails hypothesis, which together with the example descriptions entails classificaitns

Example: Gary Larson sketch where Thad the caveman shows his mates how to use a stick to grill a lizard

Relevance-Based Learning

points out relevant features

functional dependencies

generalization that derives the value of one predicate from another one

learning method

deductive

does not create any new knowledge

main effect

limitation of the hypothesis space

allows deductive generalizations from single examples

Examples: inference from nationality to spoken language

Inductive Logic Programming

learning based on predicate logic

learning methods

discovery of new predicates and new knowledge

extensions of decision trees to predicate logic

main effects

reduction of the hypothesis space to include only theories that are consistent with prior knowledge

smaller hypotheses by using prior knowledge to formulate new rules

Example: Learning of family relationships

generating the difinition of `Grandparent` becomes much easier if `Parent` is available
too complex for decision-tree learning

Learning in Multi-Agent Systems

as opposed to learning in individual agents

principal categories

differentiating features

learning coordination

learning with other agents

Principal Categories

of learning in multi-agent systems

centralized learning

one agent performs all relevant activities

no interaction required

may include multiple agents with the same learning goals

decentralized learning

several agents are engaged in the same learning process

requires interaction and collaboration between agents

Differentiating Features

for multi-agent learning approaches

degree of decentralization
distributed, parallelized

interaction-specific features
level, persistence, frequency, patterns, variability of interaction

involvement-specific features
relevance of involvement
roles played by agents

goal-specific features
type of improvement intended by learning
compatibility of learning goals across different agents

learning method
rote learning, learning by instruction, example, practice, analogy, discovery

learning feedback

supervised learning

reinforcement learning

unsupervised learning

Credit Assignment

Whose fault was it?

inter-agent credit assignment

several agents are involved in a learning activity that results in a performance change

who is responsible for the change?

intra-agent credit assignment

which of the internal components of an agent involved in learning is responsible for a performance change?

general question

what action carried out by which agent contributed to what extent to the performance change?

related question

what knowledge, what inferences, and what decisions led to an action

Learning and Activity Coordination

improving the performance of the overall system

reinforcement learning

agents try to maximize the amount of reinforcement they receive from the environment or from an instructor

isolated reinforcement learning

individual agents use reinforcement learning to achieve their own goals

no communication or collaboration about the learning processes or results

collaborative reinforcement learning

agents communicate to decide on individual and group actions

agents have some insight into each others learning processes, and share some results

Learning with Other Agents

organizational roles

assignments of roles in teams in order to learn more effectively

environmental conditions

adaptation to changes in the environment
mutual exchange of pertinent information
especially for buyer and seller agents in electronic marketplaces

team competitions

improvements in playing against competitors
learning from more experienced agents

Explanation-Based Learning

to improve cooperative problem solving

inefficiencies in coordinated behavior

identification of underlying causes

rectification of decisions or actions

learning analysis

problem-solving traces are collected and

analyzed

explanations are generated for relevant decisions

agent models

the explanations should be generated with the

model of the agent in mind

Learning and Communication

learn more and less by exchanging information

reduced communication

learning may lead to less conversations among agents

instead of asking, agents learn themselves

reduced learning

communcation can reduce the need for learning

instead of learning, agents ask other agents

balancing learning and communication

may depend on bottlenecks in computational power or bandwidth

Summary - Learning Agents

Learning in agents

- acquisition and generation of new knowledge
- performance improvement as goal

Learning Agent Model

- components, representation, feedback, prior knowledge

Learning Methods

- inductive learning, neural networks,
- reinforcement learning, genetic algorithms

Knowledge and Learning

- explanation-based learning, relevance
- information
- inductive logic programming, maching learning

Learning in Multi-Agent Systems

- de-centralized, interaction, integration of results