Reinforcement Learning

Written by: https://www.fiverr.com/jaclinejerry
Employer URL: http://www.jevan.com.au/
Date: 28 October 2016

What is Reinforcement Learning?

Definition...

Reinforcement Learning is a type of *Machine Learning*, and thereby also a branch of *Artificial Intelligence*. It allows machines and software agents to automatically determine the ideal behaviour within a specific context, in order to maximize its performance. Simple reward feedback is required for the agent to learn its behaviour; this is known as the reinforcement signal.

There are many different algorithms that tackle this issue. As a matter of fact, Reinforcement Learning is defined by a specific type of problem, and all its solutions are classed as Reinforcement Learning algorithms. In the problem, an agent is supposed decide the best action to select based on his current state. When this step is repeated, the problem is known as a *Markov Decision Process*.

Further details... Reinforcement Learning Examples

Why Reinforcement Learning?

Motivation...

Reinforcement Learning allows the machine or software agent to learn its behaviour based on feedback from the environment. This behaviour can be learnt once and for all, or keep on adapting as time goes by. If the problem is modelled with care, some Reinforcement Learning algorithms can converge to the global optimum; this is the ideal behaviour that maximises the reward.

This automated learning scheme implies that there is little need for a human expert who knows about the domain of application. Much less time will be spent designing a solution, since there is no need for hand-crafting complex sets of rules as with *Expert Systems*, and all that is required is someone familiar with Reinforcement Learning.

Tell me more! Reinforcement Learning Discussions

How does Reinforcement Learning work?

Technology...

As mentioned, there are many different solutions to the problem. The most popular, however, allow the software agent to select an action that will maximise the reward in the long term (and not only in the immediate future). Such algorithms are know to have infinite horizon.

In practice, this is done by learning to estimate the value of a particular state. This estimate is adjusted over time by propagating part of the next state's reward. If all the states and all the actions are tried a sufficient amount of times, this will allow an optimal policy to be defined; the action which maximises the value of the next state is picked.

Show me! [Tutorials on Reinforcement Learning]

When does Reinforcement Learning fail?

Limitations...

There are many challenges in current Reinforcement Learning research. Firstly, it is often too memory expensive to store values of each state, since the problems can be pretty complex. Solving this involves looking into value approximation techniques, such as *Decision Trees* or *Neural Networks*. There are many consequence of introducing these imperfect value estimations, and research tries to minimise their impact on the quality of the solution.

Moreover, problems are also generally very modular; similar behaviours reappear often, and modularity can be introduced to avoid learning everything all over again. Hierarchical approaches are common-place for this, but doing this automatically is proving a challenge. Finally, due to limited perception, it is often impossible to fully determine the current state. This also affects the performance of the algorithm, and much work has been done to compensate this *Perceptual Aliasing*.

What about... [Reinforcement Learning Papers]

Who uses Reinforcement Learning?

Applications...

The possible applications of Reinforcement Learning are abundant, due to the genericness of the problem specification. As a matter of fact, a very large number of problems in *Artificial Intelligence* can be fundamentally mapped to a decision process. This is a distinct advantage, since the same theory can be applied to many different domain specific problem with little effort.

In practice, this ranges from controlling robotic arms to find the most efficient motor combination, to robot navigation where collision avoidance behaviour can be learnt by negative feedback from bumping into obstacles. Logic games are also well-suited to Reinforcement Learning, as they are traditionally defined as a sequence of decisions

Examples

- Robotics: Quadruped Gait Control, Ball Acquisition(Robocup)
- Control: Helicopters
- Operations Research: Pricing, Routing, Scheduling
- Game Playing: Backgammon, Solitaire, Chess, Checkers
- Human Computer Interaction: Spoken Dialogue Systems
- Economics/Finance: Trading

More still... [Applications of Reinforcement Learning] http://reinforcementlearning.ai-depot.com/

MDP Vs RL

Markov decision process

- Set of states S, set of actions A
- Transition probabilities to next states T(s, a, a0)
- Reward functions R(s)

RL is based on MDPs, but

- Transition model is not known
- Reward model is not known

MDP computes an optimal policy

RL learns an optimal policy

Types of RL

1. Passive Vs Active

Passive: Agent executes a fixed policy and evaluates it

Active: Agents updates policy as it learns

2. Model based Vs Model free

Model-based: Learn transition and reward model, use it to get optimal policy

Model free: Derive optimal policy without learning the model

In Detail Please open these link:

Here is all example and pseudo code.

http://gandalf.psych.umn.edu/users/schrater/schrater_lab/courses/AI2/rl1.pdf

Reinforcement Learning in Artificial Intelligence

Full chapter: https://webdocs.cs.ualberta.ca/~sutton/papers/barto-sutton-97.pdf

Evolutionary Algorithms for Reinforcement Learning
https://arxiv.org/pdf/1106.0221.pdf

AN APPLICATION OF REINFORCEMENT LEARNING

https://www.csustan.edu/sites/default/files/honors/documents/journals/Stirrings/Darmousseh.pdf

Reinforcement Learning and the Reward Engineering Principle
http://www.danieldewey.net/reward-engineering-principle.pdf

Reinforcement Learning A Survey

https://www.cs.cmu.edu/~tom/10701_sp11/slides/Kaelbling.pdf

Distinctive Image Features

Distinctive invariant features from images that can be used to perform reliable matching between different views of an object or scene. The features are invariant to image scale and rotation, and are shown to provide robust matching across a a substantial range of affine distortion,change in 3D viewpoint, addition of noise, and change in illumination.The features are highly distinctive, in the sense that a single feature can be correctly matched with high probability against a large database of features from many images

An approach to using these features for object recognition.The recognition proceeds by matching individual features to a database of features from known objects using a fast nearest-neighbour algorithm, followed by a Hough transform to identify clusters belonging to a single object, and finally performing verification through least-squares solution for consistent pose parameters. This approach to recognition can robustly identify objects among clutter and occlusion while achieving near real-time performance.

The cost of extracting these features is minimized by taking a cascade filtering approach,in which the more expensive operations are applied only at locations that pass an initial test.

Following are the major stages of computation used to generate the set of image features:

1. Scale-space extrema detection: The first stage of computation searches over all scales

and image locations. It is implemented efficiently by using a difference-of-Gaussian

function to identify potential interest points that are invariant to scale and orientation.

2. Keypoint localization: At each candidate location, a detailed model is fit to determine location and scale. Keypoints are selected based on measures of their stability.

3. Orientation assignment: One or more orientations are assigned to each keypoint locationbased on local image gradient directions. All future operations are performed on image data that has been transformed relative to the assigned orientation, scale, and location for each feature, thereby providing invariance to these transformations.

4. Keypoint descriptor: The local image gradients are measured at the selected scale in the region around each keypoint. These are transformed into a representation that allows for significant levels of local shape distortion and change in illumination.

Feature Descriptors, Detection and Matching

http://www.cs.toronto.edu/~kyros/courses/2503/Handouts/features.pdf

Related research
http://www.robots.ox.ac.uk/~vgg/research/affine/det_eval_files/lowe_ijcv2004.pdf

Computer Vision System Toolbox

Computer Vision System Toolbox™ provides algorithms, functions, and apps for designing and simulating computer vision and video processing systems. You can perform feature detection, extraction, and matching; object detection and tracking; motion estimation; and video processing. For 3-D computer vision, the system toolbox supports camera calibration, stereo vision, 3-D reconstruction, and 3-D point cloud processing. With machine learning based frameworks, you can train object detection, object recognition, and image retrieval systems. Algorithms are available as MATLAB® functions, System objects™, and Simulink® blocks.

For rapid prototyping and embedded system design, the system toolbox supports fixed-point arithmetic and C-code generation.

- Key Features
- Object Detection and Recognition
- Camera Calibration
- Stereo Vision
- 3-D Point Cloud Processing
- Object Tracking and Motion Estimation
- Feature Detection, Extraction, and Matching
- Code Generation and Fixed Point

  For More detail:

  https://www.mathworks.com/products/computer-vision/

Object Recognition and Tracking in Video Sequences

Object recognition in computer vision is the task of finding a given object in an image or video sequence. Object recognition is one of the hardest challenges for computer vision systems today. Humans find this task extremely trivial and can recognize objects even if they are rotated, translated, scaled or partially obstructed from view. Many approaches have been applied to object recognition in single still images or still images of the object taken from different perspectives and in different poses. We propose a supervised learning technique to extend object recognition to video sequences. The task is to be able to recognize the object and be able to track it in a video sequence. Some of the challenges posed by this particular task include recognizing the object when looking at it from a different perspective and pose, recognizing the object when it is partially occluded and tracking the object while it is in motion.

For More Detail & Components Analysis

http://www.cs.dartmouth.edu/~lorenzo/teaching/cs134/Archive/Spring2009/final/NimitGeethmala/final.html

Human Detection and Character Recognition in TV-Style Movies

In order to recognize the same person at different time instances in avideo sequence, the outward appearance of the person has to be described and learned with an appropriate model. The diversity in which humans can appear makes the task of human detection and character recognition to a particularly challenging problem. TV style movies provide an uncontrolled, realistic working environment for human detection and character recognition. Possible applications range from surveillance (e.g., intrusion detection) and security applications (e.g., person identification) to image retrieval or semiautomatic image annotation (e.g., automatic labeling of faces in personal photo albums).

More detailed overview on the different stages

https://lear.inrialpes.fr/pubs/2007/Kla07/informatiktage.pdf

Re-Identification of Visual Targets in Camera Networks

A Comparison of Techniques

Re-identification is still an open problem in computer vision. The enormous possible variations from camera to camera in illumination, pose, color or all of those combined, introduce large appearance changes on the people detected, which make the problem very difficult to overcome.

Furthermore we evaluate five different classifiers: three fixed distance metrics, one learned distance metric and a classifier based on sparse representation, novel to the field of re-identification.

Click on here;

http://vislab.isr.ist.utl.pt/wp-content/uploads/2012/12/11-iciar-figueira.pdf

Real-Time Security Application to Identify the Distance and Size of an Object with CCD Camera:

This study develops an algorithm to identify the distance and size of an object and, in this respect, empirical studies were conducted. Images obtained from two security cameras were processed and the distance and size of the object was measured.

For More Detail:

https://www.researchgate.net/publication/274312435_Real-Time_Security_Application_to_Identify_the_Distance_and_Size_of_an_Object_with_CCD_Camera

Face Recognition from Still Images to Video Sequences: A Local-Feature-Based Framework:

http://jivp.eurasipjournals.springeropen.com/articles/10.1155/2011/790598

3D Exploitation of 2D Imagery

https://www.ll.mit.edu/publications/journal/pdf/vol20_no1/20_1_8_Cho.pdf