

Sistemi Intelligenti Reinforcement Learning: Fuzzy Reinforcement Learning

Alberto Borghese
Università degli Studi di Milano
Laboratorio di Sistemi Intelligenti Applicati (AIS-Lab)
Dipartimento di Scienze dell'Informazione
borghese@di.unimi.it



A.A. 2012-2013

1/20

<http://borghese.di.unimi.it/>



Sommario



- Fuzzy RL
- Stochastic Learning Automata

A.A. 2012-2013

2/20

<http://homes.dsi.unimi.it/~borghese/>



Clever Pac-man

N.A.Borghese, A.Rossini and C.Quadri (2012) Clever Pac-man,
Proceedings of the 21st Italian Workshop on Neural Nets,
WIRN2011, Frontiers in Artificial Intelligence and Applications,
IOS Press (Apolloni, Bassis, Esposito, Morabito eds.), pp.11-19.

*Applied Intelligent Systems Laboratory
Computer Science Department
University of Milano
<http://ais-lab.dsi.unimi.it>*



Motivation



How can we make a computer agent play Pac-man?

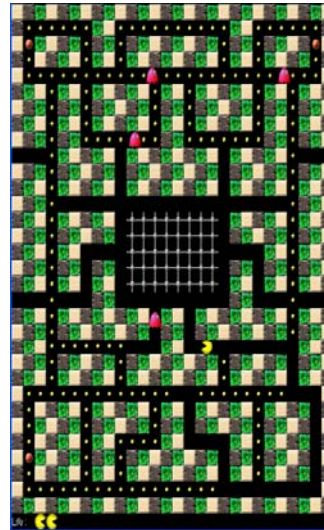


The Pac-man game



Arcade computer game

- An agent that moves in a maze. The agent is a stylized yellow mouth that opens/closes.
- The maze is constituted of corridors paved with (yellow) pills.
- When all pills are eaten the agent can move to the next game level.
- Some enemies, with the shape of pink ghosts, are present, that go after the pacman.
- Special pills, called power pills (pink spheres) are present among the pills. They allow the pacman to eat the ghosts but their effect lasts for a limited amount of time.
- Each eaten pill is worth one point, while each eaten ghost is worth 200, 400, 600, 800 points (first, second, third ghost).



A.A. 2012-2013

5/20



Pac-man as a learning agent



No a-priori information is available.

The environment (maze structure, ghosts and pills position, ghosts behavior) is not known to the pac-man
→ **environment identification.**

Large number of cells ($\cong 30 \times 32 = 960$) and situations.

Ghosts behavior has also to be specified.


Reinforcement learning is explored here. Fuzzy state definition allows managing the number of cells

Agent:

- Elements: State, Actions, Rewards, Value function.
- Policy: Action = f(State).
- Learning machinery.

Environment:

- Ghosts behavior.



A.A. 2012-2013

6/20



The ghosts

In the original game design (*Susan Lammers: "Interview with Toru Iwatani, the designer of Pac-Man", Programmers at Work 1986*), the four ghosts had different personalities:

Ghost #1, chases directly after Pac- man.

Ghost #2, positions himself a few dots in front of Pac- man mouth (if these two ghosts and the Pac-man are inside the same corridor a sandwich movement occurs).

Ghost #3 and #4, move randomly.



In the present implementation all the four ghosts can assume all three possible behaviors depending on the situation of the game (the state).

Ghosts have to escape the Pac-man when the power pill is active.

The more the game progresses the more the ghosts have to aim to the Pac-man.

A.A. 2012-2013

7/20



The ghosts behavior

At each step each ghost has to decide if moving north, south, east, west.


Shy behavior. The ghost moves away from the closest ghost. This allows distributing the ghosts inside the maze. When the power pill is active, the ghosts tend to move as far as possible from the Pac-man. *The direction the maximize the increment of distance is chosen.* When ties are present, the Pac-man makes a randomized choice to avoid stereotyped behavior.

Random behavior. It chooses an admissible direction randomly.

Hunting behavior. The ghost chooses the direction of the minimum path to the Pac-man. Minimum path has to be updated at each step as the Pac-man moves. The Floyd-Warshall algorithm is used to pre-compute the minimum path, distance between pairs of cells, for each cell of the maze, at game loading time.

Defence behavior. The ghosts go in the area in which the pills density is maximum.

To this aim the maze is subdivided into nine partially overlapped areas: $\{0 - \frac{1}{2}; \frac{1}{4} - \frac{3}{4}; \frac{1}{2} - 1\}$ and the ghost aims to the center of the area waiting for the Pac-man.



A.A. 2012-2013

8/20




The Fuzzy behavior implementation



At each step each ghost chooses among the four possible behaviors: shy, random, hunting and defence, according to a **fuzzy policy**.

Input fuzzy variables are:

- distance between the ghost and the Pac-man
- distance with the nearest ghost.
- frequency of the Pac-man eating pills.
- life time of the Pac-man (that is associated to its ability).



A set of rules have been designed like for instance:

- *If pacman_near AND skill_good, Then hunting_behavior*
- *If pacman_near AND skill_med AND pill_med, Then hunting_behavior*
- *If pacman_near AND skill_med AND pill_far, Then hunting_behavior*
- *If pacman_med AND skill_good AND pill_far, Then hunting_behavior*
- *If pacman_med AND skill_med AND pill_far, Then hunting_behavior*
- *If pacman_far AND skill_good AND pill_far, Then hunting_behavior*

Input class boundaries are chosen so that ghosts have hunting as preferred action (four times the other actions) in real game situations.

At start all ghosts are grouped in the center.

A.A. 2012-2013

9/20



The Pac-man and fuzzy Q-learning

Fuzzy description of the state is mandatory to avoid combinatorial explosion of the number of the states.

The state of the game is described by three (fuzzy) variables:

- minimum distance from the closest pill.
- minimum distance from the closest power pill.
- minimum distance from a ghost.

Three fuzzy classes for each variable -> 27 fuzzy states.



Fuzzy aggregated state	Closest ghost	Closest pill	Closest power pill
1	Low	Low	Low
2	Low	Low	Medium
3	Low	Low	High
4	Low	Medium	Low
5	Low	Medium	Medium
6	Low	Medium	High
7	Low	High	Low
8	Low	High	Medium
9	Low	High	High
10	Medium	Low	Low
11	Medium	Low	Medium
12	Medium	Low	High
13	Medium	Medium	Low
14	Medium	Medium	Medium
15	Medium	Medium	High
16	Medium	High	Low
17	Medium	High	Medium
18	Medium	High	High
19	High	Low	Low
20	High	Low	Medium
21	High	Low	High
22	High	Medium	Low
23	High	Medium	Medium
24	High	Medium	High
25	High	High	Low
26	High	High	Medium
27	High	High	High

A.A. 2012-2013

10/20



Q-learning



Actors

- State (fuzzy states) – {s}
- Actions (Go to Pill, Go to Power Pill, Avoid Ghost, Go after Ghost) – {a}
- Rewards (number of points earned in particular situations, e.g. Pill eaten), {r}

Functions


Related to environment, not known to the agent:

- Environment evolution: $s_{t+1} = g(s_t, a_t)$.
- Reward: points gained $r_{t+1} = r(s_t, a_t, s_{t+1})$

Related to the agent, optimized through learning:

- Policy: $a_t = f(s_t)$
- Value function: $Q = Q(s_t, a_t)$

$$Q(s_t, a_t) = Q(s_t, a_t) + \alpha[r_{t+1} + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)]$$



Fuzzy State of the Pac-man



We measure the state:

- The distance from the closest ghost, c_1 .
- The distance from the closest pill, c_2 .
- The distance from the closest power pill, c_3 .


Each element can fall in more than one state at each time step

We compute the membership to each fuzzy state s_j as:

$$\mu(s_j) = \frac{\sum_{i=1}^3 m(c_i)}{3}$$

With $m(\cdot)$ degree of membership of the measurement c_i to one of the fuzzy classes (small, medium, large) associated to each state variable (distance from closest ghost, closest pill, closest power pill).

More than one state can be active at each time step and the degrees of activity, $\mu(s_j)$ add to one.





Fuzzy Q-learning

The value function for the state s^* , constituted of all the states with their membership value, from which the Pac-man moves, with action a , receives contribution from all the states associated to the next position of the Pac-man inside the maze:



$$Q(s_t^*, a_t) = \frac{1}{n} \sum_{i=1}^n \mu(s_{t,i}) q(s_{t,i}, a_t)$$

where $q(\cdot)$ is updated using Q-learning strategy as:

$$q(s_t, a_t) = q(s_t, a_t) + \alpha_{s,a} [r + \gamma \cdot \max_{a'} Q(s_{t+1}, a') - q(s_t, a_t)]$$

α is chosen as:

$$\alpha_{s,a} = \frac{1}{\sum_{\tau=0}^{t-1} \mu(s_\tau(s_t, a_t))}$$

That is a natural extension of running average computation and it is inversely proportional to the cumulative membership of all the states active at that time step.



Implementation issues of Pac-man policy

$a_t = \{ \text{Go to Pill, Go to Power Pill, Avoid Ghost, Go to Ghost} \}$

Policy: $a_t = f(s_t)$



Go to Pill. The Pac-man always goes to the closest pill, independently on the position of the ghosts. If ties occur the choice is randomized to avoid stereotyped behavior.

Go to Power Pill. Similar as above.

Go to Ghost. Similar as above.

Avoid Ghost. If only the closest ghost is considered, the Pac-man would easily run into a second ghost. The move that minimizes the weighted distance with all the ghost could be considered, but this would move the Pac-man in a small area close to the corners of the maze. We have implemented a weighted distance computed only inside a small area around the actual position of the Pac-man (that changes at each time step). Moreover, in case of ties, the Pac-man chooses the direction that leads to the closest power pill (if still present in the maze).



Additional implementation issues



Few heuristics have been introduced:

Persistence (cf. DeLooze, L.L.; Viner, W.R.; "Fuzzy Q-learning in a nondeterministic environment: developing an intelligent Ms. Pac-Management", *Computational Intelligence and Games, 2009. CIG 2009*. pp.162-169, 7-10 Sept. 2009). Forcing the same action for n steps (n=5 here).



Persistence removal. When power pill effect ends. A brisk change of behavior is often observed.

Taboo. Inhibits the Pac-man to return in the previous state.



Parameters role



Rewards. The death of the Pac-man receives instant reward of -1000. A less negative reward was not enough to compensate all the positive points earned during a typical game. A more negative reward made the Pac-man "depressed" and little inclined to look for pills.

Fuzzy classes boundary: $d=5$, $d=12$ and $d = 25$ were assumed as maximum distance for the classes: low, medium and large. These values have been experimentally set analyzing the game results.

Pills reward: no particular effect was observed when the value was in the range $[0.1 \div 1]$.



Greediness of the policy



Greediness of the policy: ϵ -greedy policy is fundamental to obtain very good results.

With random policy (blue) little points are gained.


Some more points can be gained if the Pac-man always chooses “avoid ghosts” unless he has eaten the Power Pill (orange).

Maximum reward is obtained when Q-learning with ϵ -greedy policy with $\epsilon=0.1$ choice is adopted and $r = 0.1$ per pill (yellow).

A high reward is obtained when Q-learning with ϵ -greedy policy with $\epsilon=0.1$ choice is adopted and $r = 1$ per pill (green).

Less reward is obtained with Q-learning with greedy policy (brown).

An even small reward is obtained with Q-learning with greedy policy, when fuzzy classes boundaries are different: $d = \{6, 18, 30\}$ (cyan).



A.A. 2012-2013

17/20



Conclusion and further developments



- Highest score was around 4,500 and reported in *DeLooze, L.L.; Viner, W.R.; "Fuzzy Q-learning in a nondeterministic environment: developing an intelligent Ms. Pac-Man agent", Computational Intelligence and Games, 2009. CIG 2009, pp.162-169, 7-10 Sept. 2009.* We obtain here a large improvement in the score.
- Fuzzy approach has made RL approach feasible.

- We have only considered the bonus represented by power pills.
- A single scheme was used.
- Fuzzy classes boundaries were not optimized.
- A human player elaborates strategies both in chasing and escaping that are based on a global “view” of the game. This would require a much elaborate learning machinery than “simple” RL.

- Here is the Pac-man learning live....

A.A. 2012-2013

18/20



Launch Fuzzy Pac-man



- Spostarsi nella cartella *bin dell'applicazione*.
- Lanciare il file main:
`java pacman.PacmanMAIN`



Sommario



- Fuzzy RL
- Stochastic Learning Automata