

Reinforcement Learning



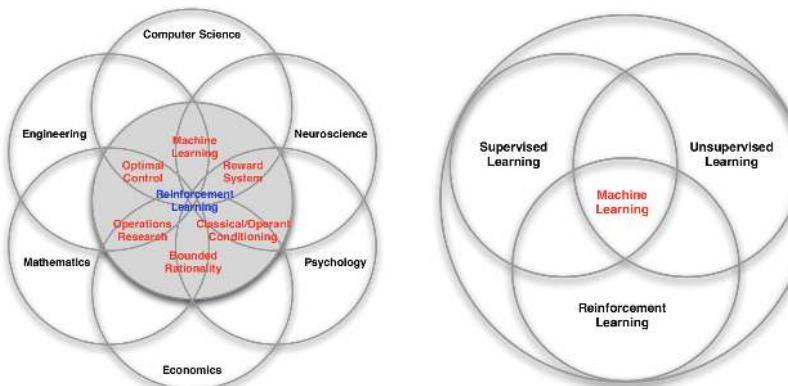
Jean-Paul Comet (projet Bioinfo Formelle)

EPU dept GB-BIMB - 5ème année
Université Côte d'Azur

18 décembre 2023

inspired from David Silver, Antoine
Cornuéjols, Sutton and Barto, 1998
<http://webdocs.cs.ualberta.ca/~sutton/book/the-book.html...>

Introduction : general schema



1 Introduction

2 Notions d'utilité et de politique

3 ϵ -greedy Algorithm

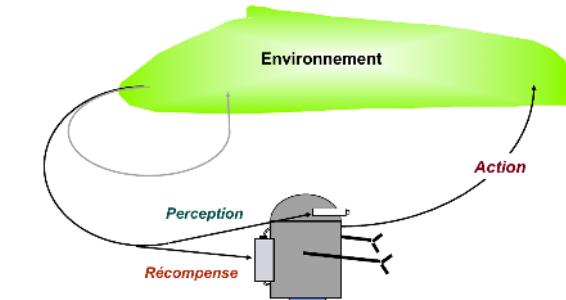
4 Markov Decision Processes : MDP

5 Quand le mode est connu

6 Quand le monde est connu (et monde fini)

7 Apprentissage avec généralisation

Introduction : general schema



Characteristics of Reinforcement Learning

What makes reinforcement learning different from other machine learning paradigms ?

- There is no supervisor, only a reward signal
- Feedback is delayed, not instantaneous
- Time really matters (sequential, non i.i.d data)
- Agent's actions affect the subsequent data it receives

Introduction : general schema

Exploration vs. Exploitation Dilemma

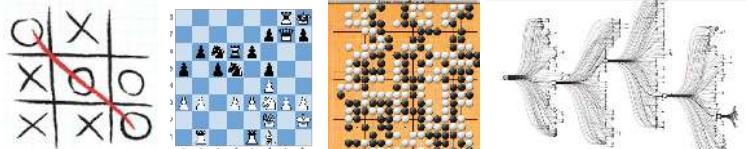
- Online decision-making involves a fundamental choice :
 - Exploitation** Make the best decision given current information
 - Exploration** Gather more information
- The best long-term strategy may involve short-term sacrifices
- Gather enough information to make the best overall decisions

Examples

- Restaurant Selection
 - Exploitation** Go to your favourite restaurant
 - Exploration** Try a new restaurant
- Online Banner Advertisements
 - Exploitation** Show the most successful advert
 - Exploration** Show a different advert
- Oil Drilling
 - Exploitation** Drill at the best known location
 - Exploration** Drill at a new location
- Game Playing
 - Exploitation** Play the move you believe is best
 - Exploration** Play an experimental move

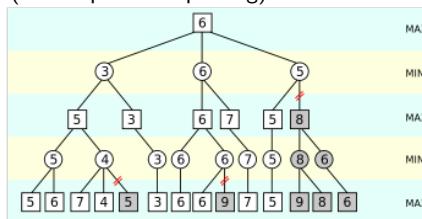
Introduction : general schema

- Application to conventional Games : All aspects of the states are fully observable



- Technique : MinMax (with Alpha-Beta pruning) Effective for

- Modest branching factor
- When a good evaluation function is available

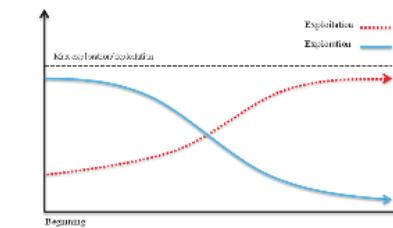
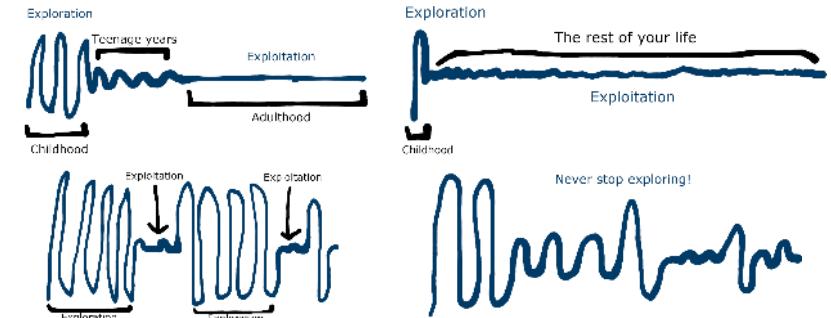


Go : branching factor 250 [Chess : 31] + no good eval. function

- Generalisation

- Often, no "endgame", so no **exploration** to the leaves possible
A reinforcement signal from time to time
- We don't play against a MIN type opponent. So no worst case.
We want to maximize the expectation of gain in an unknown environment (at the beginning)

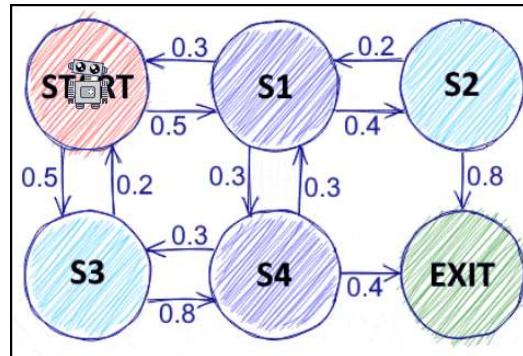
Introduction : general schema



Introduction : Basic notations

- Discrete Time : t
- States : $s_t \in S$
- Actions : $a_t \in \mathcal{A}(s_t)$
- Rewards : $r_t \in \mathcal{R}(s_t)$
- The agent : $s_t \rightarrow a_t$
- The environment : $(s_t, a_t) \rightarrow s_{t+1}, r_{t+1}$
- Policy : $\pi_t : S \rightarrow \mathcal{A}$ T (transitions) + R (rewards)
with $\pi_t(s, a) = \text{probability of } a_t = a \text{ when } s_t = s$
- Transitions and Rewards depend only on current state and action :
Markovian process

Markov Decision Processes : MDP



- ① Given any state and action, s and a , the *transition probability* of each possible next state, s' , is

$$\mathcal{P}_{ss'}^a = \Pr\{s_{t+1} = s' \mid s_t = s, a_t = a\}$$
- ② Given any current state and action, s and a , together with any next state, s' , the expected value of the next reward is

$$\mathcal{R}_s^a = E\{r_{t+1} \mid s_t = s, a_t = a, s_{t+1} = s'\}$$
- ③ One can also define the reward dynamics in terms of the expected next reward given just the current state and action, that is, by

$$\mathcal{R}_s^a = E\{r_{t+1} \mid s_t = s, a_t = a\}$$
. This quantity is related to our as follows :

$$\mathcal{R}_s^a = \sum_{s'} \mathcal{P}_{ss'}^a \mathcal{R}_{s'}^a$$

Value functions : recursive relationships

$$\begin{aligned} V^\pi(s) &= E_\pi\{\mathcal{R}_t \mid s_t = s\} = E_\pi\left\{\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \mid s_t = s\right\} \\ &= E_\pi\left\{r_{t+1} + \gamma \sum_{k=0}^{\infty} \gamma^k r_{t+k+2} \mid s_t = s\right\} \\ &= \sum_a \pi(s, a) \sum_{s'} \mathcal{P}_{ss'}^a \left[\mathcal{R}_{s'}^a + \gamma E_\pi\left\{\sum_{k=0}^{\infty} \gamma^k r_{t+k+2} \mid s_t = s'\right\} \right] \\ &= \sum_a \pi(s, a) \sum_{s'} \mathcal{P}_{ss'}^a [\mathcal{R}_{s'}^a + \gamma V^\pi(s')] \end{aligned}$$

We also have :

$$\begin{aligned} Q^\pi(s, a) &= E_\pi\{\mathcal{R}_t \mid s_t = s, a_t = a\} \\ &= \sum_{s'} \mathcal{P}_{ss'}^a [\mathcal{R}_{s'}^a + \gamma V^\pi(s')] \end{aligned}$$

Value functions

Value functions : estimate how good it is for the agent to be in a given state

- "how good" : defined in terms of future rewards that can be expected, or, to be precise, in terms of expected return.
- A policy, π , is a mapping from each state, $s \in \mathcal{S}$, and action, $a \in \mathcal{A}$, to the probability $\pi(s, a)$ of taking action a when in state s .

state-value function for policy π

Value of a state s under a policy π , denoted $V^\pi(s)$, is the expected return when starting in s and following π thereafter. For MDPs, we can define formally as

$$V^\pi(s) = E_\pi\{\mathcal{R}_t \mid s_t = s\} = E_\pi\left\{\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \mid s_t = s\right\}$$

action-value function for policy π

the value of taking action a in state s under a policy π , denoted $Q^\pi(s, a)$ is the expected return starting from s , taking the action a , and thereafter following policy π :

$$Q^\pi(s, a) = E_\pi\{\mathcal{R}_t \mid s_t = s, a_t = a\} = E_\pi\left\{\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \mid s_t = s, a_t = a\right\}$$

Optimal value function $V(s)$ & optimal policy

- Une politique est une application $\pi : S \rightarrow A$
- Valeur optimale d'un état :

$$V^*(s) = \max_\pi V_\pi(s) = \max_\pi E_\pi\left(\sum_{t=0}^{+\infty} \gamma^t r^t\right)$$

- La fonction de valeur optimale V^* est unique

$$V^*(s) = \max_a \sum_{s'} \mathcal{P}_{ss'}^a [\mathcal{R}_{s'}^a + \gamma V^*(s')]$$

- Une politique stationnaire optimale existe :

$$\pi^*(s) = a^* = \operatorname{ArgMax}_a \sum_{s'} \mathcal{P}_{ss'}^a [\mathcal{R}_{s'}^a + \gamma V^*(s')]$$

(à connaître : $\mathcal{P}_{ss'}^a$ and $\mathcal{R}_{s'}^a$)

- Valeur optimale d'une action dans un état :

$$\begin{aligned} Q^*(s, a) &= \max_{\pi} Q_{\pi}(s, a) = E_{\pi} \{ r_{t+1} + \gamma V^*(s_{t+1}) \mid s_t = s, a_t = a \} \\ &= \sum_{s'} \mathcal{P}_{ss'}^a [R_{ss'}^a + \gamma \max_{a'} Q^*(s', a')] \end{aligned}$$

- Théorème :

$\pi^*(s) = \operatorname{ArgMax}_a Q^*(s, a)$ est une politique optimale

(rien à connaître)

Prog. dynamique : Évaluation de politique (Algo V)

Input : π , the policy to be evaluated

Algorithm parameter :

a small threshold $\theta > 0$ determining accuracy of estimation

Initialize an array $V(s)$, for all $s \in S^+$, arbitrarily

Repeat :

$\Delta \leftarrow 0$

For each $s \in S$:

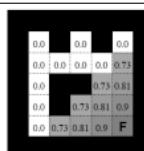
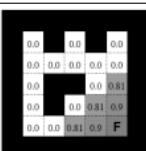
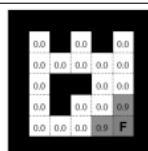
$v \leftarrow V(s)$

$V(s) \leftarrow \sum_a \pi(s, a) \sum_{s'} \mathcal{P}_{ss'}^a [R_{ss'}^a + \gamma V(s')]$

$\Delta \leftarrow \max(\Delta, |v - V(s)|)$

until $\Delta < \theta$ (a small positive number)

Output $V \sim v^{\pi}$



Algorithme d'itération de valeur

- Évaluation de politique :** Pour une politique donnée π , calculer la fonction d'utilité d'état $V_{\pi}(s)$:
State-value function for policy π :-

$$V^{\pi}(s) = E_{\pi} \{ R_t \mid s_t = s \} = E_{\pi} \left\{ \sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \mid s_t = s \right\}$$

Bellman equation for V_{π} :

$$V^{\pi}(s) = \sum_a \pi(s, a) \sum_{s'} \mathcal{P}_{ss'}^a [R_{ss'}^a + \gamma V^{\pi}(s')]$$

⇒ a system of $|S|$ simultaneous linear equations

- Évaluation itérative d'une politique (prog. dyn.)**

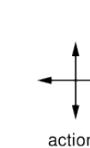
Principe : l'équation de point fixe de Bellman peut fournir une procédure itérative d'approximation successive de la fonction d'utilité V_{π} .

$$V_{k+1}(s) \leftarrow \sum_a \pi(s, a) \sum_{s'} \mathcal{P}_{ss'}^a [R_{ss'}^a + \gamma V_k(s')]$$

Propagation :

$$V_0 \rightarrow V_1 \rightarrow \dots \rightarrow V_k \rightarrow V_{k+1} \rightarrow \dots \rightarrow V^{\pi}$$

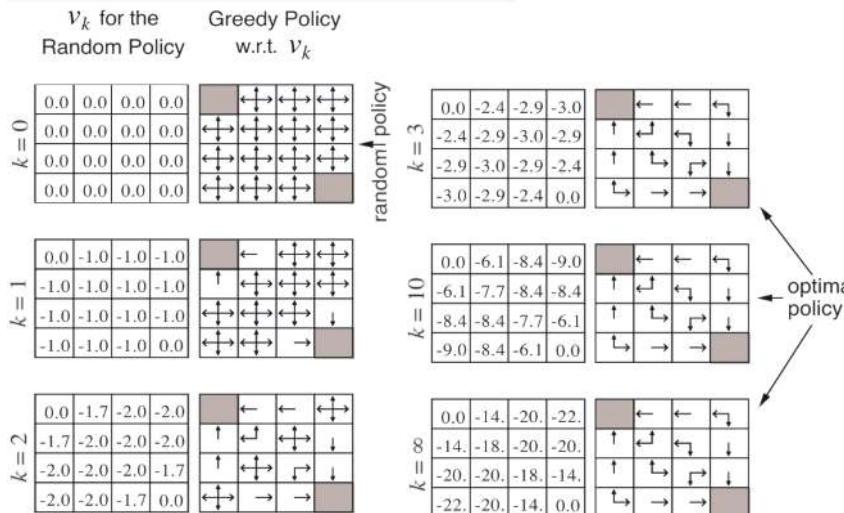
DP : Evaluation d'une politique : exemple



	1	2	3
4	5	6	7
8	9	10	11
12	13	14	

$r = -1$
on all transitions

DP : Evaluation d'une politique : exemple



Comment améliorer une politique ?

Lorsqu'on change la valeur V , la politique n'est peut être pas la meilleure.

Avec les nouvelles valeurs de V , la meilleure politique est :

$$\begin{aligned}\pi'(s) &= \underset{a}{\operatorname{ArgMax}} Q^\pi(s, a) \\ &= \underset{a}{\operatorname{ArgMax}} \sum_{s'} \mathcal{P}_{ss'}^a [\mathcal{R}_{ss'}^a + \gamma V^\pi(s')]\end{aligned}$$

Prog. dynamique : Évaluation de politique (Algo Q)

Input : π , the policy to be evaluated

Algorithm parameter :

a small threshold $\theta > 0$ determining accuracy of estimation

Initialize an array $Q(s, a)$, for all $s \in S^+$, $a \in A$ arbitrarily

Repeat :

$$\Delta \leftarrow 0$$

For each $s \in S$:

For each $a \in A$:

$$q \leftarrow Q(s, a)$$

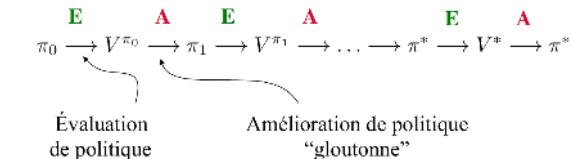
$$Q(s, a) \leftarrow R(s, a) + \gamma \sum_{s'} T(s, a, s') \max_{a'} Q(s', a')$$

$$\Delta \leftarrow \max(\Delta, |q - Q(s, a)|)$$

until $\Delta < \theta$

Output $Q \sim Q^\pi$

Itération de politique



Initialisation arbitraire de π

Faire

calcul de la fonction de valeur avec π

$$V_\pi(s) = R(s, \pi(s)) + \gamma \sum_{s' \in S} T(s, \pi(s), s') V_\pi(s')$$

Amélioration de la politique à chaque état

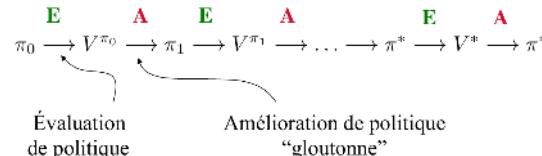
$$\pi'(s) \leftarrow \max_a \left(R(s, a) + \gamma \sum_{s' \in S} T(s, a, s') V_\pi(s') \right)$$

$$\pi' := \pi$$

Jusqu'à ce qu'aucune amélioration ne soit possible

Garantie de convergence vers une politique optimale

Itération de politique



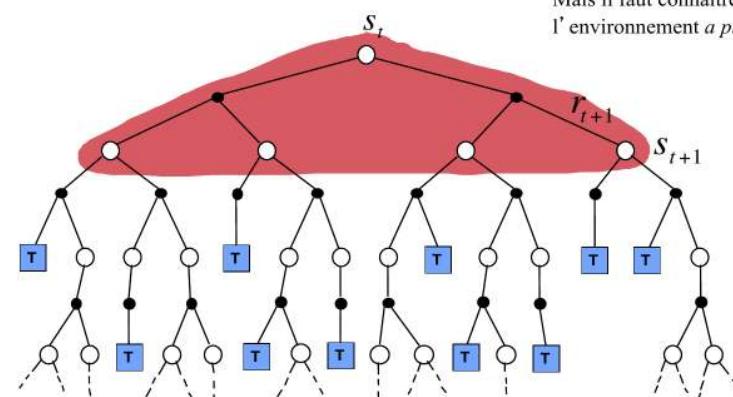
- Initialization**
 $V(s) \in \mathbb{R}$ and $\pi(s) \in \mathcal{A}(s)$ arbitrarily for all $s \in \mathcal{S}$
- Policy Evaluation**
Loop:
 $\Delta \leftarrow 0$
Loop for each $s \in \mathcal{S}$:
 $v \leftarrow V(s)$
 $V(s) \leftarrow \sum_{s',r} p(s',r|s,\pi(s)) [r + \gamma V(s')]$
 $\Delta \leftarrow \max(\Delta, |v - V(s)|)$
until $\Delta < \theta$ (a small positive number determining the accuracy of estimation)
- Policy Improvement**
 $\text{policy-stable} \leftarrow \text{true}$
For each $s \in \mathcal{S}$:
 $\text{old-action} \leftarrow \pi(s)$
 $\pi(s) \leftarrow \arg\max_a \sum_{s',r} p(s',r|s,a) [r + \gamma V(s')]$
If $\text{old-action} \neq \pi(s)$, then $\text{policy-stable} \leftarrow \text{false}$
If policy-stable , then stop and return $V \approx v_*$ and $\pi \approx \pi_*$; else go to 2

Le modèle du monde est inconnu

TD learning : cf. Dynamic Programming

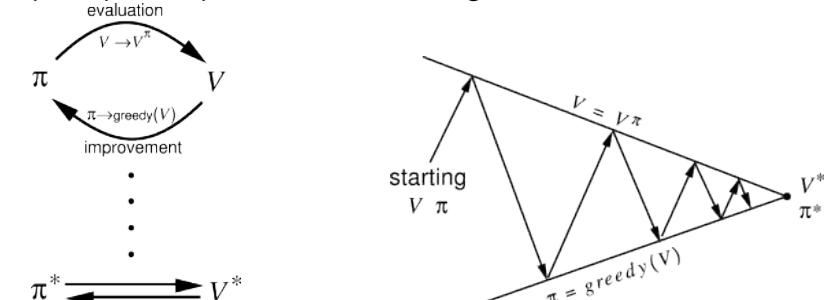
$$V(s_t) \leftarrow E_\pi \{ r_{t+1} + \gamma V(s_{t+1}) \}$$

On calcule l'espérance.
Mais il faut connaître l'environnement *a priori*.



Itération généralisée de politique

Generalized Policy Iteration (GPI) : Toute interaction d'étape d'évaluation de politique et d'étape d'amélioration de politique indépendamment de leur granularité



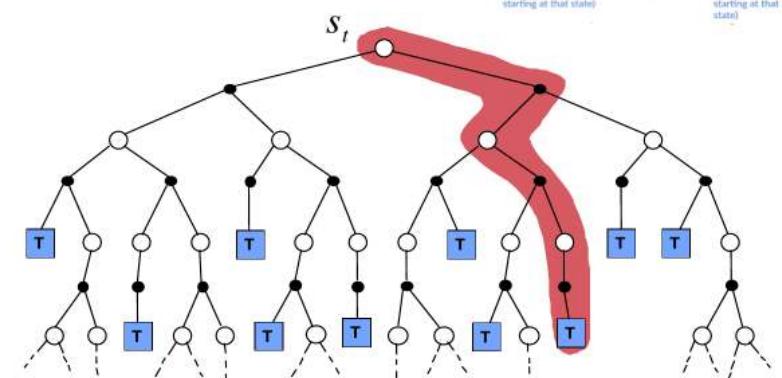
On peut se permettre aussi de ne pas calculer à chaque étape les nouvelles valeurs de V pour tous les états. Nombreux travaux, nombreuses variantes...

Le modèle du monde est inconnu

Monte Carlo method

$$V(S_t) \leftarrow V(S_t) + \alpha [G_t - V(S_t)]$$

New value of state t
Former estimation of value of state t (= Expected return starting at that state)
Learning Rate
Return of timestep
Former estimation of value of state t (= Expected return starting at that state)



e.g., the TD(0) backup :

$$V(s_t) \leftarrow V(s_t) + \alpha[r_{t+1} + \gamma V(s_{t+1}) - V(s_t)]$$

As a training example :

$$\{ \text{description of } \underbrace{s_t}_{\text{input}}, \underbrace{r_{t+1} + \gamma V(s_{t+1})}_{\text{target output}} \}$$

- En principe, oui :
 - Réseaux de neurones artificiels
 - Arbres de décision
 - Méthodes de régression multivariées
 - etc.
- Mais l'App. par R. a des exigences particulières :
 - Apprendre tout en agissant
 - S'adapter à des mondes non stationnaires