

Course guide

270420 - APRNS - Unsupervised and Reinforcement Learning

Last modified: 10/07/2024

Unit in charge: Barcelona School of Informatics
Teaching unit: 723 - CS - Department of Computer Science.

Degree: BACHELOR'S DEGREE IN ARTIFICIAL INTELLIGENCE (Syllabus 2021). (Compulsory subject).

Academic year: 2024 **ECTS Credits:** 6.0 **Languages:** Catalan, Spanish

LECTURER

Coordinating lecturer: JAVIER BÉJAR ALONSO - MARIO MARTÍN MUÑOZ

Others: Primer quadrimestre:
JAVIER BÉJAR ALONSO - 11, 12
MARIO MARTÍN MUÑOZ - 11, 12

PRIOR SKILLS

Basic knowledge of Deep Learning and Machine Learning.

DEGREE COMPETENCES TO WHICH THE SUBJECT CONTRIBUTES

Specific:

CE18. To acquire and develop computational learning techniques and to design and implement applications and systems that use them, including those dedicated to the automatic extraction of information and knowledge from large volumes of data.

Generical:

CG2. To use the fundamental knowledge and solid work methodologies acquired during the studies to adapt to the new technological scenarios of the future.

CG4. Reasoning, analyzing reality and designing algorithms and formulations that model it. To identify problems and construct valid algorithmic or mathematical solutions, eventually new, integrating the necessary multidisciplinary knowledge, evaluating different alternatives with a critical spirit, justifying the decisions taken, interpreting and synthesizing the results in the context of the application domain and establishing methodological generalizations based on specific applications.

Transversal:

CT6. Autonomous Learning. Detect deficiencies in one's own knowledge and overcome them through critical reflection and the choice of the best action to extend this knowledge.

Basic:

CB5. That the students have developed those learning skills necessary to undertake later studies with a high degree of autonomy

TEACHING METHODOLOGY

The classes are divided into theory, problem and laboratory sessions.

In the theory sessions, knowledge of the subject will be developed, interspersed with the presentation of new theoretical material with examples and interaction with the students in order to discuss the concepts.

In the laboratory classes, small practices will be developed using tools and using specific libraries that will allow you to practice and reinforce the knowledge of the theory classes.

LEARNING OBJECTIVES OF THE SUBJECT

- 1.To distinguish the kind of problems can be modeled as a reinforcement learning problem and identify the techniques that can be applied to solve them.
- 2.To understand the need, fundamentals, and particularities of behavior learning and the differences it has from supervised and unsupervised machine learning.
- 3.To understand the most important algorithms and state of the art in the area of learning by reinforcement
- 4.To know how to computationally formalize a real world problem as learning by reinforcement and know how to implement in the most current environments the learning algorithms that solve them
- 5.Know the problems that can be modeled with deep unsupervised algorithms
- 6.Understand the particularities of deep unsupervised algorithms
- 7.Know the most important algorithms and the state of the art of deep unsupervised learning
- 8.Knowing how to implement and apply deep learning algorithms to a problem using the most current environment

STUDY LOAD

Type	Hours	Percentage
Hours small group	30,0	20.00
Hours large group	30,0	20.00
Self study	90,0	60.00

Total learning time: 150 h

CONTENTS

Introduction: Behavior Learning in Agents and description of main elements in Reinforcement Learning

Description:

Intuition, motivation and definition of the reinforcement learning (RL) framework. Key elements in RL.

Finding optimal policies using Dynamic Programming

Description:

How to learn the optimal policy with full knowledge of the world model: algebraic solution, policy iteration and value iteration.

Introduction to Model-Free approaches.

Description:

Basic algorithms for reinforcement learning: Monte-Carlo, Q-learning, Sarsa, TD(λ). The need for Exploration. Differences between On-policy and Off-policy methods.

Function approximation in Reinforcement Learning

Description:

Need for function approximation and Incremental methods in RL. The Gradient Descent approach. RL with Linear function approximation. The deadly triad for function approximation in RL. Batch methods and Neural Networks for function Approximation.

Deep Reinforcement Learning (DRL)

Description:

Introducing Deep Learning in RL. Dealing with the deadly triad with the DQN algorithm. Application to the Atari games case. Evolutions of the DQN algorithm: Double DQN, Prioritized Experience Replay, multi-step learning and Distributional value functions. Rainbow: the state-of-the-art algorithm in discrete action space.

Policy gradient methods

Description:

What to do in continuous action spaces. How probabilistic policies allow to apply the gradient method directly in the policy network. The REINFORCE algorithm. The Actor-Critic algorithms. State-of-the-art algorithms in continuous action spaces: DDPG, TD3 and SAC.

Advanced Topics: How to deal with sparse rewards

Description:

The problem of the sparse reward. Introduction to advanced exploration techniques: curiosity and empowerment in RL. Introduction to curriculum learning to ease the learning of the goal. Hierarchical RL to learn complex tasks. The learning of Universal Value Functions and Hindsight Experience Replay (HER).

Reinforcement Learning in the multi-agent framework

Description:

Learning of behaviors in environment where several agents act. Learning of cooperative behaviors, Learning of competitive behaviors, and mixed cases. State-of-the-art algorithms. The special case of games: The AlfaGo case and the extension to AlfaZero.

Introduction: Deep unsupervised learning

Description:

Introduction to the need for deep unsupervised learning and its applications

Autoregressive models

Description:

Introduction to learning probability distributions defined as autoregressive distributions and main models

Normalizing flows

Description:

Introduction to normalized flows for learning probability distributions

Latent variables models

Description:

Introduction to models based on latent variables and variational autoencoders



Generative Adversarial Networks

Description:

Introduction to generative adversarial networks, conditional and unconditional generation, attribute disentanglement

Denosing Diffusion networks

Description:

Introduction to models based on noise diffusion, denoising networks, conditioning, multimodal generation

Self supervised learning

Description:

Introduction to self-supervised learning for feature-generating networks and embeddings, contrastive and non contrastive methods, masking

ACTIVITIES

Introduction: Behavior Learning in Agents and description of main elements in Reinforcement Learning

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Finding optimal policies using Dynamic Programming

Description:

How to learn the optimal policy with full knowledge of the world model: algebraic solution, policy iteration and value iteration.

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Introduction to Model-Free approaches. Monte-Carlo, Q-learning, Sarsa, TD(λ)

Description:

Development of the corresponding topic

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Function approximation in RL

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Deep Reinforcement Learning (DRL)

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Policy gradient methods

Description:

What to do in continuous action spaces. How probabilistic policies allow to apply the gradient method directly in the policy network. The REINFORCE algorithm. The Actor-Critic algorithms. State-of-the-art algorithms in continuous action spaces: DDPG, TD3 and SAC.

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Advanced Topics: How to deal with sparse rewards

Description:

The problem of the sparse reward. Introduction to advanced exploration techniques: curiosity and empowerment in RL. Introduction to curriculum learning to ease the learning of the goal. Hierarchical RL to learn complex tasks. The learning of Universal Value Functions and Hindsight Experience Replay (HER).

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Reinforcement Learning in the multi-agent framework

Description:

Learning of behaviors in environment where several agents act. Learning of cooperative behaviors, Learning of competitive behaviors, and mixed cases. State-of-the art algorithms. The special case of games: The AlfaGo case and the extension to AlfaZero.

Full-or-part-time: 13h

Self study: 9h

Theory classes: 2h

Laboratory classes: 2h

Control of the reinforcement learning part

Specific objectives:

1, 2, 3, 4

Related competencies :

CE18. To acquire and develop computational learning techniques and to design and implement applications and systems that use them, including those dedicated to the automatic extraction of information and knowledge from large volumes of data.

CG2. To use the fundamental knowledge and solid work methodologies acquired during the studies to adapt to the new technological scenarios of the future.

CG4. Reasoning, analyzing reality and designing algorithms and formulations that model it. To identify problems and construct valid algorithmic or mathematical solutions, eventually new, integrating the necessary multidisciplinary knowledge, evaluating different alternatives with a critical spirit, justifying the decisions taken, interpreting and synthesizing the results in the context of the application domain and establishing methodological generalizations based on specific applications.

CT6. Autonomous Learning. Detect deficiencies in one's own knowledge and overcome them through critical reflection and the choice of the best action to extend this knowledge.

Full-or-part-time: 2h

Guided activities: 2h

Introduction: Deep unsupervised learning

Description:

Introduction to the need for deep unsupervised learning and its applications

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Autoregressive models

Description:

Introduction to learning probability distributions defined as autoregressive distributions and main models

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Normalizing flows

Description:

Introduction to normalized flows for learning probability distributions

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Latent variables models

Description:

Introduction to models based on latent variables and variational autoencoders

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Generative Adversarial Networks

Description:

Introduction to generative adversarial networks, conditional and unconditional generation, attribute disentanglement

Full-or-part-time: 10h

Self study: 6h

Theory classes: 2h

Laboratory classes: 2h

Denosing Diffusion networks and Self supervised learning

Description:

Introduction to models based on noise diffusion, denoising networks, conditioning, multimodal generation

Full-or-part-time: 15h

Self study: 9h

Theory classes: 2h

Laboratory classes: 4h

Unsupervised learning syllabus control

Specific objectives:

5, 6, 7, 8

Related competencies :

CB5. That the students have developed those learning skills necessary to undertake later studies with a high degree of autonomy
CE18. To acquire and develop computational learning techniques and to design and implement applications and systems that use them, including those dedicated to the automatic extraction of information and knowledge from large volumes of data.

CG2. To use the fundamental knowledge and solid work methodologies acquired during the studies to adapt to the new technological scenarios of the future.

CG4. Reasoning, analyzing reality and designing algorithms and formulations that model it. To identify problems and construct valid algorithmic or mathematical solutions, eventually new, integrating the necessary multidisciplinary knowledge, evaluating different alternatives with a critical spirit, justifying the decisions taken, interpreting and synthesizing the results in the context of the application domain and establishing methodological generalizations based on specific applications.

CT6. Autonomous Learning. Detect deficiencies in one's own knowledge and overcome them through critical reflection and the choice of the best action to extend this knowledge.



GRADING SYSTEM

The subject will include the following assessment acts:

- Reports of the laboratory activities, which must be delivered within the deadline indicated for each session (roughly, 2 weeks). Based on a weighted average of the grades of these reports, a laboratory grade will be calculated, L.
- A first partial exam, taken towards the middle of the course, of the material seen until then. Let P1 be the grade obtained in this exam.
- On the designated day within the exam period, a second partial exam of the subject not covered by the first partial. Let P2 be the grade obtained in this exam.

The three grades L, P1, P2 are between 0 and 10.

The final grade of the subject will be: $0.4 \cdot L + 0.3 \cdot P1 + 0.3 \cdot P2$

BIBLIOGRAPHY

Basic:

- Sutton, Richard S; Barto, Andrew G. Reinforcement learning : an introduction. Second edition. Cambridge, Massachusetts: The MIT Press, [2020]. ISBN 9780262039246.
- Morales, Miguel. Grokking deep reinforcement learning. Shelter Island: Manning Publications, 2020. ISBN 9781617295454.
- Foster, D. Generative deep learning: teaching machines to paint, write, compose, and play. 2nd ed. Sebastopol: O'Reilly Media, Incorporated, 2023. ISBN 9781098134143.
- Cheong, S.Y. Hands-on image generation with TensorFlow: a practical guide to generating images and videos using deep learning. Birmingham: Packt Publishing, 2020. ISBN 9781838821104.

Complementary:

- Zai, Alexander; Brown, Brandon. Deep reinforcement learning in action. Shelter Island, NY: Manning Publications Co, 2020. ISBN 9781617295430.
- Babcock, J.; Bali, R. Generative AI with Python and TensorFlow 2: harness the power of generative models to create images, text, and music. Birmingham, England ; Mumbai: Packt Publishing, 2021. ISBN 9781800208506.