

Année universitaire 2024/2025

Intelligence Artificielle, Systèmes, Données - 2e année de Master

Crédits ECTS : 60

LES OBJECTIFS DE LA FORMATION

Cette formation d'excellence offre de solides connaissances en mathématiques appliquées et conception de systèmes d'intelligence artificielle afin de couvrir l'ensemble des problématiques de traitement de données massives que rencontre les entreprises. Elle met l'accent sur l'articulation entre apprentissage automatique, gestion et fouille de grandes masses de données, paradigmes du Big Data, représentation des connaissances, le traitement des données et sur les méthodologies récemment développées.

Les objectifs de la formation :

- Former des informaticiens capables de maîtriser les problèmes conceptuels, sémantiques et algorithmique soulevés par l'intelligence artificielle et la science des données
- Développer une compréhension générale et en profondeur des différentes facettes de l'IA
- Renforcer des connaissances théoriques des étudiants et leur transmettre une expérience pratique de l'Intelligence Artificielle et des Sciences des Données

Lieux des cours :

Master IASD initial : à compter de janvier 2025, les cours seront dispensés au 16 bis rue de l'Estrapade, 75005 Paris.

Master IASD en alternance : les cours sont dispensés sur le site de [PariSanté Campus](#).

PRÉ-REQUIS OBLIGATOIRES

- Titulaires d'un diplôme BAC+4 (240 crédits ECTS) ou équivalent à Dauphine, d'une université ou d'un autre établissement de l'enseignement supérieur dans les domaines suivants : informatique, mathématiques appliquées avec un attrait pour l'informatique et l'algorithme
- Etudiants en dernière année d'école d'ingénieur (ou ayant obtenu un diplôme d'ingénieur) en lien avec les thématiques de la formation
- Le master s'adresse aux étudiants qui se destinent à des carrières de recherche et de développement dans le domaine de l'intelligence artificielle et des sciences des données. Le Master recrute parmi les meilleurs étudiants ayant validé une première année de Master en informatique, ou en mathématiques appliquées avec un attrait pour l'informatique et l'algorithmique.
- Notez, qu'un goût particulier pour l'IA et les sciences des données est absolument indispensable !

POURSUITE D'ÉTUDES

Carrière académique, doctorat (Université, CNRS, INRIA, CEA, CNES, INRA, ect...) et carrière R&D (Google, Facebook, Criteo, Keyrus, Amazon, 1000mercis, IBM, HAVAS, AXA, BNP Parisbas,...)

PROGRAMME DE LA FORMATION

- Semestre 3
 - UE obligatoires S3

- Foundations of Machine Learning
- Optimization for Machine Learning
- Data acquisition, extraction and storage
- Data Science Lab
- Deep learning for image analysis
- Reinforcement learning
- Large language models
- Semestre 4
 - UE optionnelles S4
 - Advanced machine learning
 - Incremental learning, game theory and applications
 - Point Clouds and 3D Modelling
 - Knowledge graphs, description logics, reasoning on data
 - Graph analytics
 - Machine learning on Big Data
 - Computational social choice
 - Monte-Carlo search and games
 - Deep reinforcement learning et applications
 - Privacy for Machine Learning
 - No SQL databases
 - Non-convex inverse problems
 - Mathematics of deep learning
 - Science des données au Collège de France
 - Bloc stage
 - Stage

DESCRIPTION DE CHAQUE ENSEIGNEMENT

Advanced machine learning

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

This research-oriented module will focus on advanced machine learning algorithms, in particular in the Bayesian setting

1) Bayesian Machine Learning (with Moez Draief, chief data scientist CapGemini)

- Bayesian linear regression
- Gaussian Processes (i.e. kernelized Bayesian linear regression)
- Approximate Bayesian Inference
- Latent Dirichlet Allocation

2) Bayesian Deep Learning (with Julyan Arbel, CR INRIA)

- MCMC methods
- variationnal methods

3) Advanced Recommendation Techniques (with Clement Calauzene, Criteo)

Compétence à acquérir :

Probabilistic, Bayesian ML and recommendation systems

Mode de contrôle des connaissances :

- Chaque étudiant devra présenter un papiers de recherche

Computational social choice

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

The aim of this course is to give an overview of the problems, techniques and applications of computational social choice, a multidisciplinary topic at the crossing point of computer science (especially artificial intelligence, operations research, theoretical computer science, multi-agent systems, computational logic, web science) and economics.

The course consists of the analysis of problems arising from the aggregation of preferences of a group of agents from a computational perspective. On the one hand, it is concerned with the application of techniques developed in computer science, such as complexity analysis or algorithm design, to the study of social choice mechanisms, such as voting procedures or fair division algorithms. On the other hand, computational social choice is concerned with importing concepts from social choice theory into computing.

The course will focus on normative aspects, computational aspects, and real-world applications (including some case studies).

Program:

1. Introduction to social choice and computational social choice.
2. Preference aggregation, Arrow's theorem and how to escape it.
3. Voting rules: informational basis and normative aspects.
4. Voting rules : computation. Voting on combinatorial domains.
5. Strategic issues: strategyproofness, Gibbard and Satterthwaite's theorem, computational resistance to manipulation, other forms of strategic behaviour.
6. Multiwinner elections. Public decision making and participatory budgeting.
7. Communication issues in voting: voting with incomplete preferences, elicitation protocols, communication complexity, low-communication social choice.
8. Fair division.
9. Matching under preferences.
10. Specific applications and case studies (varying every year): rent division, kidney exchange, school assignment, group recommendation systems...

Compétence à acquérir :

N/S

Mode de contrôle des connaissances :

Written exam by default.

Bibliographie, lectures recommandées :

References:

* Handbook of Computational Social Choice (F. Brandt, V. Conitzer, U. Endriss, J. Lang, A. Procaccia, eds.), Cambridge University Press, 2016. Available for free online.

* Trends in Computational Social Choice (U. Endriss, ed), 2017. Available for free online.

Data Science Lab

ECTS : 4

Volume horaire : 27

Description du contenu de l'enseignement :

Students enrolled in this class will form groups and choose one topic among a list of proposed topics in the core areas of the master such as supervised or unsupervised learning, recommendation, game AI, distributed or parallel data-science, etc. The topics will generally consist in applying a well-established technique on a novel data-science challenge or in applying recent research results on a classical data-science challenge. Either way, each topic will come with its own novel scientific challenge to address. At the end of the module, the students will give an oral presentation to demonstrate their methodology and their findings. Strong scientific rigor as well as very good engineering and communication skills will be necessary to complete this module successfully.

Compétence à acquérir :

The goal of this module is to provide students with a hands-on experience on a novel data-science/AI challenge using state-of-the-art tools and techniques discussed during other classes of this master.

Data acquisition, extraction and storage

ECTS : 4

Volume horaire : 27

Description du contenu de l'enseignement :

The objective of this course is to present the principles and techniques used to acquire, extract, integrate, clean, preprocess, store, and query datasets, that may then be used as input data to train various artificial intelligence models. The course will consist on a mix of lectures and practical sessions. We will cover the following aspects:

- Web data acquisition (Web crawling, Web APIs, open data, legal issues)
- Information extraction from semi-structured data
- Data cleaning and data deduplication
- Data formats and data models
- Storing and processing data in databases, in main memory, or in plain files
- Introduction to large-scale data processing with MapReduce and Spark
- Ontology-based data access

Compétence à acquérir :

Understanding:

- how to acquire data from a variety of sources and in a variety of formats
- how to extract structured data from unstructured or semi-structured data
- how to format, integrate, clean data sets
- how to store and access data sets

Mode de contrôle des connaissances :

Project (50% of the grade) and in-class written assessment (50% of the grade)

Deep learning for image analysis

ECTS : 4

Volume horaire : 27

Description du contenu de l'enseignement :

Deep learning has achieved formidable results in the image analysis field in recent years, in many cases exceeding human performance. This success opens paths for new applications, entrepreneurship and research, while making the field very competitive.

This course aims at providing the students with the theoretical and practical basis for understanding and using deep learning for image analysis applications.

Program to be followed

The course will be composed of lectures and practical sessions. Moreover, experts from industry will present practical applications of deep learning.

Lectures will include:

- Artificial neural networks, back-propagation algorithm
- Convolutional neural networks
- Design and optimization of a neural architecture
- Analysis of neural network function
- Image classification and segmentation
- Auto-encoders and generative networks
- Transformers
- Current research trends and perspectives

During the practical sessions, the students will code in Python, using Keras or Pytorch. They will be confronted with the practical problems linked to deep learning: architecture design; optimization schemes and hyper-parameter selection; analysis of results.

Compétence à acquérir :

Deep learning for image analysis: theoretical foundations and applications

Mode de contrôle des connaissances :

Deep reinforcement learning et applications

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

What you will learn in this class?

- Intro and Course Overview
- Supervised Learning behaviors
- Principles of Reinforcement Learning
- Policy Gradients
- Actor-Critic Algorithms (A2C, A3C and Soft AC)
- Value Function Methods
- Deep RL with Q-functions
- Advanced Policy Gradient (DDPG, Twin Delayed DDPG)
- Trust Region & Proximal Policy Optimization (TRPO, PPO)
- Optimal Control and Planning
- Model-Based Reinforcement Learning
- Model-Based Policy Learning
- Exploration and Stochastic Bandit in RL
- Exploration with Curiosity and Imagination
- Offline RL and Generalization issues
- Offline RL and Policy constraints

Why you should choose this course about DRL?

- DRL Is a very promising type of learning as it does not need to know the solution
- DRL Only needs the rules and good rewards
- DRL Combines the best aspects of deep learning and reinforcement learning.
- DRL has achieved impressive results in games, robotic, finance and many more fields

References

- Bertsekas, Dynamic Programming and Optimal Control, Vols I and II
- Goodfellow, Bengio, Deep Learning
- Powell, Approximate Dynamic Programming
- Puterman, Markov Decision Processes: Discrete Stochastic Dynamic Programming
- Sutton & Barto, Reinforcement Learning: An Introduction
- Szepesvari, Algorithms for Reinforcement Learning

Compétence à acquérir :

WHAT YOU WILL ACQUIRE IN THIS CLASS?

- Understand principles of Deep Reinforcement Learning (DRL)
- Know main DRL algorithms
- Get some intuition about what DRL is good and not good at?
- Program DRL algorithms

Foundations of Machine Learning

ECTS : 4

Volume horaire : 27

Description du contenu de l'enseignement :

The course will introduce the theoretical foundations of machine learning, review the most successful algorithms with their

theoretical guarantees, and discuss their application in real world problems. The covered topics are:

- **Part 1:** Supervised Learning Theory: the batch setting
 - Intro
 - Surrogate Losses
 - Uniform Convergence and PAC Learning
 - Empirical Risk Minimization and ill-posed problems
 - Concentration Inequalities
 - Universal consistency, PAC Learnability
 - VC dimension
 - Rademacher complexity
 - Non Uniform Learning and Model Selection
 - bias-variance tradeoff
 - Structural Minimization Principle and Minimum Description Length Principle
 - Regularization
- **Part 2:** Supervised Learning Theory and Algorithms in the Online Setting
 - Foundations of Online Learning
 - Beyond the Perceptron algorithm
- **Partie 3:** Ensemble Methods and Kernels Methods
 - SVMs, Kernels
 - Kernel approximation algorithms in the primal
 - Ensemble methods: bagging, boosting, gradient boosting, random forests
- **Partie 4:** Algorithms for Unsupervised Learning
 - Dimensionality reduction: PCA, ICA, Kernel PCA, ISOMAP, LLE
 - Representation Learning
 - Expectation Maximization, Latent models and Variational methods

Compétence à acquérir :

The aim of this course is to provide the students with the fundamental concepts and tools for developing and analyzing machine learning algorithms.

Mode de contrôle des connaissances :

- Each student will have to have the role of scribe during one lecture, taking notes during the class and sending the notes to the teacher in pdf.
- Final exam

Bibliographie, lectures recommandées :

The most important book:

- Shalev-Shwartz, S., & Ben-David, S. (2014). Understanding machine learning: From theory to algorithms. Cambridge university press.

Also:

- Mohri, M., Rostamizadeh, A., & Talwalkar, A. (2012). Foundations of machine learning. MIT press.

- Vapnik, V. (2013). The nature of statistical learning theory. Springer science & business media.

- Bishop Ch. (2006). Pattern recognition and machine learning. Springer

- Friedman, J., Hastie, T., & Tibshirani, R. (2001). The elements of statistical learning (Vol. 1, No. 10). New York, NY, USA:: Springer series in statistics.

- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). An introduction to statistical learning (Vol. 112). New York: springer.

Graph analytics

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

The objective of this course is to give students an overview of the field of graph analytics. Since graphs form a complex and expressive data type, we need methods for representing graphs in databases, manipulating, querying, analyzing and mining them. Moreover, graph applications are very diverse and need specific algorithms.

The course presents new ways to model, store, retrieve, mine and analyze graph-structured data and some examples of applications.

Lab sessions are included allowing students to practice graph analytics: modeling a problem into a graph database and

performing analytical tasks over the graph in a scalable manner.

Program

- Graph analytics
 - Networks properties and models
 - Link Analysis : PageRank and its variants
 - Community detection
- Frameworks for parallel graph analytics
 - Pregel – a model for parallel-graph computing
 - GraphX Spark – unifying graph- and data –parallel computing
- Machine learning with graphs
- Applications : process mining and analysis

Practical work : graph analytics with GraphX and Neo4J

Compétence à acquérir :

Modeling a problem into a graph model and performing analytical tasks over the graph in a scalable manner.

Bibliographie, lectures recommandées :

References

Ian Robinson, Jim Weber, Emil Eifrem, Graph Databases, Editeur : O'Reilly (4 juin 2013), ISBN-10: 1449356265

Eric Redmond, Jim R. Wilson, Seven Databases in Seven Weeks - A Guide to Modern Databases and the NoSQL Movement, Publisher: Pragmatic Bookshelf

Grzegorz Malewicz, Matthew H. Austern, Aart J.C Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. 2010. Pregel: a system for large-scale graph processing, SIGMOD '10, ACM, New York, NY, USA, 135-146

Xin, Reynold & Crankshaw, Daniel & Dave, Ankur & Gonzalez, Joseph & J. Franklin, Michael & Stoica, Ion. (2014). GraphX: Unifying Data-Parallel and Graph-Parallel Analytics.

Michael S. Malak and Robin East, Spark GraphX in Action, Manning, June 2016

Incremental learning, game theory and applications

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

This course will focus on the behavior of learning algorithms when several agents are competing against one another: specifically, what happens when an agent that follows an online learning algorithm interacts with another agent doing the same? The natural language to frame such questions is that of game theory, and the course will begin with a short introduction to the topic, such as normal form games (in particular zero-sum, potential, and stable games), solution concepts (such as dominated/rationalizable strategies, Nash, correlated and coarse equilibrium notions, ESS), and some extensions (Blackwell approachability). Subsequently, we will examine the long-term behavior of a wide variety of online learning algorithms (fictitious play, regret-matching, multiplicative/exponential weights, mirror descent and its variants, etc.), and we will discuss applications to generative adversarial networks (GANs), traffic routing, prediction, and online auctions.

[1] Nicolò Cesa-Bianchi and Gábor Lugosi, Prediction, learning, and games, Cambridge University Press, 2006.

[2] Drew Fudenberg and David K. Levine, The theory of learning in games, Economic learning and social evolution, vol. 2, MIT Press, Cambridge, MA, 1998.

[3] Sergiu Hart and Andreu Mas-Colell, Simple adaptive strategies: from regret matching to uncoupled dynamics, World Scientific Series in Economic Theory - Volume 4, World Scientific Publishing, 2013.

[4] Vianney Perchet, Approachability, regret and calibration: implications and equivalences, Journal of Dynamics and Games 1 (2014), no. 2, 181–254.

[5] Shai Shalev-Shwartz, Online learning and online convex optimization, Foundations and Trends in Machine Learning 4 (2011), no. 2, 107–194.

Compétence à acquérir :

Learning procedures when several agents are playing against one-other

Knowledge graphs, description logics, reasoning on data

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

Introduction to Knowledge Graphs, Description Logics and Reasoning on Data.

Knowledge graphs are a flexible tool to represent knowledge about the real world. After presenting some of the existing knowledge graphs (such as DBPedia, Wikidata or Yago) , we focus on their interaction with semantics, which is formalized through the use of so-called ontologies. We then present some central logical formalism used to express ontologies, such as Description Logics and Existential Rules. A large part of the course will be devoted to study the associated reasoning tasks, with a particular focus on querying a knowledge graph through an ontology. Both theoretical aspects (such as the tradeoff between the expressivity of the ontology language versus the complexity of the reasoning tasks) and practical ones (efficient algorithms) will be considered.

Program:

1. Knowledge Graphs (history and uses)
2. Ontology Languages (Description Logics, Existential Rules)
3. Reasoning Tasks (Consistency, classification, Ontological Query Answering)
4. Ontological Query Answering (Forward and backward chaining, Decidability and complexity, Algorithms, Advanced Topics)

References:

- The description logic handbook: theory, implementation, and applications. Baader et al., Cambridge University Press
- Foundations of Semantic Web Technologies, Hitzler et al., Chapman&Hall/CRC
- Web Data Management, Abiteboul et al., Cambridge University Press

Prerequisites:

- first-order logic;
- complexity (Turing machines, classical complexity classes) is a plus.

Large language models

ECTS : 4

Volume horaire : 27

Description du contenu de l'enseignement :

The course focuses on modern and statistical approaches to NLP.

Natural language processing (NLP) is today present in some many applications because people communicate most everything in language : post on social media, web search, advertisement, emails and SMS, customer service exchange, language translation, etc. While NLP heavily relies on machine learning approaches and the use of large corpora, the peculiarities and diversity of language data imply dedicated models to efficiently process linguistic information and the underlying computational properties of natural languages.

Moreover, NLP is a fast evolving domain, in which cutting-edge research can nowadays be introduced in large scale applications in a couple of years.

The course focuses on modern and statistical approaches to NLP: using large corpora, statistical models for acquisition, disambiguation, parsing, understanding and translation. An important part will be dedicated to deep-learning models for NLP.

- Introduction to NLP, the main tasks, issues and peculiarities

- Sequence tagging: models and applications
- Computational Semantics
- Syntax and Parsing
- Deep Learning for NLP: introduction and basics
- Deep Learning for NLP: advanced architectures
- Deep Learning for NLP: Machine translation, a case study

Compétence à acquérir :

- Skills in Natural Language Processing using deep-learning
- Understand new architectures

Bibliographie, lectures recommandées :

References

- Costa-jussà, M. R., Allauzen, A., Barrault, L., Cho, K., & Schwenk, H. (2017). Introduction to the special issue on deep learning approaches for machine translation. *Computer Speech & Language*, 46, 367-373.
- Dan Jurafsky and James H. Martin. *Speech and Language Processing* (3rd ed. draft): <https://web.stanford.edu/~jurafsky/slp3/>
- Yoav Goldberg. *A Primer on Neural Network Models for Natural Language Processing*: <http://u.cs.biu.ac.il/~yogo/nnlp.pdf>
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*: <http://www.deeplearningbook.org/>

Machine learning on Big Data

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

This course focuses on the typical, fundamental aspects that need to be dealt with in the design of machine learning algorithms that can be executed in a distributed fashion, typically on Hadoop clusters, in order to deal with big data sets, by taking into account scalability and robustness.

Nowadays there is an ever increasing demand of machine learning algorithms that scales over massive data sets. In this context, this course focuses on the typical, fundamental aspects that need to be dealt with in the design of machine learning algorithms that can be executed in a distributed fashion, typically on Hadoop clusters, in order to deal with big data sets, by taking into account scalability and robustness. So the course will first focus on a bunch of main-stream, sequential machine learning algorithms, by taking then into account the following crucial and complex aspects. The first one is the re-design of algorithms by relying on programming paradigms for distribution and parallelism based on map-reduce (e.g., Spark, Flink, ...). The second aspect is experimental analysis of the map-reduce based implementation of designed algorithms in order to test their scalability and precision. The third aspect concerns the study and application of optimisation techniques in order to overcome lack of scalability and to improve execution time of designed algorithm.

The attention will be on machine learning technique for dimension reduction, clustering and classification, whose underlying implementation techniques are transversal and find application in a wide range of several other machine learning algorithms. For some of the studied algorithms, the course will present techniques for a from-scratch map-reduce implementation, while for other algorithms packages like Spark ML will be used and end-to-end pipelines will be designed. In both cases algorithms will be analysed and optimised on real life data sets, by relying on a local Hadoop cluster, as well as on a cluster on the Amazon WS cloud.

References:

- Mining of Massive Datasets
<http://www.mmds.org>
- High Performance Spark - Best Practices for Scaling and Optimizing Apache Spark
Holden Karau, Rachel Warren
O'Reilly

Mathematics of deep learning

ECTS : 3

Volume horaire : 24

Monte-Carlo search and games

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

Introduction to Monte Carlo for computer games.

Monte Carlo Search has revolutionized computer games. It works well with Deep Learning so as to create systems that have superhuman performances in games such as Go, Chess, Hex or Shogi. It is also appropriate to address difficult optimization problems. In this course we will present different Monte Carlo search algorithms such as UCT, GRAVE, Nested Monte Carlo and Payout Policy Adaptation. We will also see how to combine Monte Carlo Search and Deep Learning. The validation of the course is a project involving a game or an optimization problem.

La recherche Monte-Carlo a révolutionné la programmation des jeux. Elle se combine bien avec le Deep Learning pour créer des systèmes qui jouent mieux que les meilleurs joueurs humains à des jeux comme le Go, les Echecs, le Hex ou le Shogi. Elle permet aussi d'approcher des problèmes d'optimisation difficiles. Dans ce cours nous traiterons des différents algorithmes de recherche Monte-Carlo comme UCT, GRAVE ou le Monte-Carlo imbriqué et l'apprentissage de politique de payouts. Nous verrons aussi comment combiner recherche Monte-Carlo et apprentissage profond. Le cours sera validé par un projet portant sur un jeu ou un problème d'optimisation difficile.

Bibliographie, lectures recommandées :

Bibliographie :

Intelligence Artificielle Une Approche Ludique, Tristan Cazenave, Editions Ellipses, 2011.

No SQL databases

ECTS : 3

Volume horaire : 24

Non-convex inverse problems

ECTS : 3

Volume horaire : 18

Optimization for Machine Learning

ECTS : 6

Volume horaire : 48

Description du contenu de l'enseignement :

This course delves into the mathematical underpinnings and algorithmic strategies essential for understanding and applying Machine Learning techniques. Central to the course is the exploration of optimization, a pivotal element in contemporary advancements in machine learning. This exploration encompasses fundamental approaches such as linear regression, SVMs, and kernel methods, and extends to the dynamic realm of deep learning. Deep learning has become a leading methodology for addressing a variety of challenges in areas like imaging, vision, and natural language processing. The course content is structured to provide a comprehensive overview of the mathematical foundations, algorithmic methods, and a variety of modern applications utilizing diverse optimization techniques. Participants will engage in both traditional lectures and practical numerical sessions using Python. The curriculum is divided into three parts: The first focuses on smooth and convex optimization techniques, including gradient descent and duality. The second part delves into advanced methods like non-smooth optimization and proximal methods. Lastly, the third part addresses large-scale methods such as stochastic gradient descent and automatic differentiation, with a special focus on their applications in neural networks, including both shallow and deep architectures.

Detailed Syllabus:

1. Foundational Concepts in Differential Calculus and Gradient Descent:

- Introduction to differential calculus
- Principles of gradient descent
- Application of gradient descent in optimization

2. Automatic Differentiation and Its Applications:

- Understanding the mechanics of automatic differentiation
- Implementing automatic differentiation using modern Python frameworks

3. Advanced Gradient Descent Techniques:

- In-depth study of gradient descent theory
- Accelerated gradient methods
- Stochastic gradient algorithms and their applications

4. Exploring Convex and Non-Convex Optimization:

- Fundamentals of convex analysis
- Strategies and challenges in non-convex optimization

5. Special Topics in Optimization:

- Introduction to non-smooth optimization methods
- Study of semidefinite programming (SDP)
- Exploring interior points and proximal methods

6. Large-Scale Optimization Methods and Neural Networks:

- Techniques in large-scale methods, focusing on stochastic gradient descent
- Applications of automatic differentiation in neural networks
- Overview of neural network architectures: shallow and deep networks

Bibliographie, lectures recommandées :

- Theory and algorithms: Convex Optimization, Boyd and Vandenberghe
- Introduction to matrix numerical analysis and optimization, Philippe Ciarlet
- Proximal algorithms, N. Parikh and S. Boyd
- Introduction to Nonlinear Optimization - Theory, Algorithms and Applications, Amir Beck
- Numerics: Python and Jupyter installation: use only Python 3 with Anaconda distribution.
- The Numerical Tours of Signal Processing, Gabriel Peyré
- Scikitlearn tutorial, Fabian Pedregosa, Jake VanderPlas

Point Clouds and 3D Modelling

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

Ce cours donne un panorama des concepts et techniques d'acquisition, de traitement et de visualisation des nuages de points 3D, et de leurs fondements mathématiques et algorithmiques.

Le cours aborde notamment les thèmes suivants :

Systèmes de perception 3D

Traitements et opérateurs

Recalage

Segmentation de nuages de points

Reconstruction de courbes et surfaces

Modélisation par primitives

Rendu de nuages de points et maillages

La plupart des séances sont complétées d'un TP.

Les cours sont en français, les sujets des TP sont en anglais.

Compétence à acquérir :

Compétences théoriques et pratiques sur la production, le traitement et les applications des nuages de points 3D.

Mode de contrôle des connaissances :

Comptes-rendus de TP.

Projet sur article.

Privacy for Machine Learning

ECTS : 3

Volume horaire : 24

Description du contenu de l'enseignement :

- Motivations, traditional approaches, randomized response
- Definition and properties of differential privacy
- Mechanisms for discrete/categorical data
- Mechanisms for continuous data
- Alternative notions of differential privacy
- Differential privacy for statistical learning
- Attacks and connections with robustness
- Local differential privacy and federated learning

Compétence à acquérir :

This course covers the basics of Differential Privacy (DP), a framework that has become, in the last ten years, a de facto standard for enforcing user privacy in data processing pipelines. DP methods seek to reach a proper trade-off between protecting the characteristics of individuals and guaranteeing that the outcomes of the data analysis stays meaningful.

The first part of the course is devoted the basic notion of epsilon-DP and understanding the trade-off between privacy and accuracy, both from the empirical and statistical points of view. The second half of the course will cover more advanced aspects, including the different variants of DP and the their use to allow for privacy-preserving training of large and/or distributed machine learning models.

Mode de contrôle des connaissances :

- Individual homework (Python)
- Group project on a research paper (with report and defense)

Bibliographie, lectures recommandées :

- [The Algorithmic Foundations of Differential Privacy](#), C. Dwork & A. Roth, Foundations and Trends in Theoretical Computer Science (2014)
- [Programming Differential Privacy](#), J. P. Near & C. Abueh, online book (2021)

Reinforcement learning

ECTS : 4

Volume horaire : 27

Description du contenu de l'enseignement :

- Models: Markov decision processes (MDP), multiarmed bandits and other models
- Planning: finite and infinite horizon problems, the value function, Bellman equations, dynamic programming, value and policy iteration
- Basic learning tools: Monte Carlo methods, temporal-difference learning, policy gradient
- Probabilistic and statistical tools for RL: Bayesian approach, relative entropy and hypothesis testing, concentration inequalities
- Optimal exploration in multiarmed bandits: the explore vs exploit tradeoff, lower bounds, the UCB algorithm, Thompson sampling
- Extensions: Contextual bandits, optimal exploration for MDP

Compétence à acquérir :

Reinforcement Learning (RL) refers to scenarios where the learning algorithm operates in closed-loop, simultaneously using past data to adjust its decisions and taking actions that will influence future observations. Algorithms based on RL concepts are now commonly used in programmatic marketing on the web, robotics or in computer game playing. All models for RL share a common concern that in order to attain one's long-term optimality goals, it is necessary to reach a proper balance between exploration (discovery of yet uncertain behaviors) and exploitation (focusing on the actions that have produced the most relevant results so far).

The methods used in RL draw ideas from control, statistics and machine learning. This introductory course will provide the main methodological building blocks of RL, focussing on probabilistic methods in the case where both the set of possible actions and the state space of the system are finite. Some basic notions in probability theory are required to follow the course. The course will imply some work on simple implementations of the algorithms, assuming familiarity with Python.

Mode de contrôle des connaissances :

- Individual homework (in Python)
- Final exam

Bibliographie, lectures recommandées :

Bibliographie, lectures recommandées

- M. Puterman. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons, 1994.
- R. Sutton and A. Barto. Introduction to Reinforcement Learning. MIT Press, 1998.
- C. Szepesvari. Algorithms for Reinforcement Learning. Morgan & Claypool Publishers, 2010.
- T. Lattimore and C. Szepesvari. Bandit Algorithms. Cambridge University Press. 2019.

Science des données au Collège de France

ECTS : 3

Stage

ECTS : 10

Université Paris Dauphine - PSL - Place du Maréchal de Lattre de Tassigny - 75775 PARIS Cedex 16 - 12/07/2025