

**ACADEMIC AFFAIRS OFFICE
INDIAN INSTITUTE OF TECHNOLOGY ROORKEE**

No. Acd./2360 /IAPC-114

Dated: December 21, 2021

Head, Mehta Family School of Data Science & Artificial Intelligence

The IAPC in its 114th meeting held on 01.12.2021 vide Item No. 114.2.6 considered and approved the revised syllabus of PEC i.e., AID-554: Reinforcement Learning with minor modifications.

The modified syllabus is attached as **Appendix-A**.


Assistant Registrar (Curriculum)

Encl: as above

Copy to (through e mail):-

1. All faculty
2. Head of all Departments / Centres
3. Dean, Academic Affairs
4. Associate Dean of Academic Affairs (Curriculum)
5. Channel i/ Acad portal/ Academic webpage of iitr.ac.in

INDIAN INSTITUTE OF TECHNOLOGY ROORKEE

NAME OF DEPARTMENT/CENTRE: Mehta Family School of Data Science and Artificial Intelligence

1. **Subject Code:** AID-554 **Course Title:** Reinforcement Learning
2. **Contact Hours:** **L:** 3 **T:** 1 **P:** 0
3. **Examination Duration (Hrs.):** **Theory:** 3 **Practical:** 0
4. **Relative Weightage:** **CWS:** 20-35 **PRS:** 0 **MTE:** 20-30 **ETE:** 40-50 **PRE:** 0
5. **Credits:** 4 **6. Semester:** Both **7. Subject Area:** PEC
8. **Pre-requisite:** Nil
9. **Objective:** This course aims to understand several reinforcement learning algorithms and their applications, along with emerging research trends.

10. Details of the Course

S.No.	Contents	Contact hours
1.	Basics of probability and linear algebra, Definition of a stochastic multi-armed bandit, Definition of regret, Achieving sublinear regret, UCB algorithm, KL-UCB, Thompson Sampling.	6
2.	Markov Decision Problem, policy, and value function, Reward models (infinite discounted, total, finite horizon, and average), Episodic & continuing tasks, Bellman's optimality operator, and Value iteration & policy iteration	8
3.	The Reinforcement Learning problem, prediction and control problems, Model-based algorithm, Monte Carlo methods for prediction, and Online implementation of Monte Carlo policy evaluation	8
4.	Bootstrapping; TD(0) algorithm; Convergence of Monte Carlo and batch TD(0) algorithms; Model-free control: Q-learning, Sarsa, Expected Sarsa.	6
5.	n-step returns; TD(λ) algorithm; Need for generalization in practice; Linear function approximation and geometric view; Linear TD(λ).	6
6.	Tile coding; Control with function approximation; Policy search; Policy gradient methods; Experience replay; Fitted Q Iteration; Case studies.	8
Total		42

11. Suggested Books:

S.No.	Name of Authors/Book/Publisher	Year of Publication / Reprint
1.	Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction. MIT press	2020
2.	Sugiyama, Masashi. Statistical reinforcement learning: modern machine learning approaches. CRC Press	2015
3.	Lattimore, T., and C. Szepesvári. Bandit algorithms. Cambridge University Press.	2020
4.	Boris Belousov, Hany Abdulsamad, Pascal Klink, Simone Parisi, and Jan Peters Reinforcement Learning Algorithms: Analysis and Applications. Springer International Publishing.	2021
5.	Alexander Zai and Brandon Brown, Deep Reinforcement Learning in Action. Manning press.	2020