

Proposal for a New Course

Department of intelligent systems
Indian Institute of Technology, Kanpur

1. Course No: ISXXX

2. Course Title: **Safe Learning and Control under Uncertainty**

3. **Per Week Lectures:** 3_(L), Tutorial: 0 (T), Laboratory: 0 (P), Additional Hours [0-2]: 0 (A),
Credits (L+T+P+A): 9

4. Duration of Course: 26 Lectures

5. Proposing Department/IDP: Dept. of Intelligent Systems

Other Departments/IDPs which may be interested in the proposed course:

Electrical Engineering, Aerospace Engineering, Computer Science Engineering,
Mechanical Engineering, Mathematics

Other faculty members interested in teaching the proposed course:

6. Proposing Instructor(s): Dr. Ashwin Verma

7. Course Description:

A) Objectives:

This course develops mathematical methods and computational algorithms for learning control policies subject to safety requirements. It connects constrained Markov decision processes with robust reachability, control barrier functions, and optimization-based safety filters. Students will learn to formulate safety specifications, distinguish expected constraint satisfaction from trajectory safety, and incorporate model uncertainty into controller design. The course also introduces conformal prediction and conformal risk control for statistical calibration of uncertainty and risk. Analytical exercises and computational case studies will develop students' ability to evaluate the assumptions, guarantees, and performance limitations of safe learning methods.

B) Contents (*preferably in the form of 5 to 10 broad titles*):

S. No	Broad Title	Topics	No. of Lectures
1	Dynamical systems and safety	State-space models, control-affine systems, stochastic dynamics, bounded disturbances and parametric uncertainty, state and input constraints, probabilistic and robust safety specifications.	2
2	RL foundations	Markov Decision Process (MDP), value functions and Bellman equations, dynamic programming, value and policy iteration, policy-gradient methods, REINFORCE and variance reduction.	6

3	Constrained RL	Constrained MDP, occupancy measures, linear-programming formulations, randomized policies, Lagrangian relaxation and duality, primal-dual policy optimization, constraint satisfaction in reinforcement learning.	4
4	Reachability and invariance	Backward reachability, reach-avoid problems, robust predecessor operators, controlled invariant sets and viability kernels, dynamic games (brief), introduction to HJB and HJI reachability, numerical reachability methods.	3
5	Safety-critical control	Lyapunov stability and control Lyapunov functions, forward invariance, Control Barrier Functions (CBF), CBF-QP formulations, robust CBF, input-constrained feasibility, discrete-time barrier functions, sampled-data safety.	5
6	Learning with safety filter	Safety filtering and action projection, filtered MDPs, policy optimization with safety filters, action aliasing, safe exploration, safety-performance tradeoffs.	3
7	Conformal Methods	Split conformal prediction, exchangeability and prediction coverage, conformal risk control, finite-sample risk guarantees, trajectory uncertainty quantification, distribution shift and statistical safety guarantees.	3

Total Lectures per week: 2

C) Pre-requisites, if any

1. Modern control theory, including state-space models and stability analysis.
2. Knowledge in Probability and Statistics along with MATLAB simulation experience is required.
3. Familiarity with basic constrained optimization is expected; preparatory material will be provided for students without this background.

D) Short summary for including in the Courses of Study Booklet: This course develops mathematical and computational methods for learning safe control policies under uncertainty, integrating constrained decision-making, reachability, control barrier functions, safety filters, and conformal methods for uncertainty and risk calibration.

8. References

1. Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. 2nd ed., MIT Press, 2018.
2. Dongsheng Ding, Kaiqing Zhang, Jiali Duan, Tamer Başar, and Mihailo R. Jovanovic. “Convergence and Sample Complexity of Natural Policy Gradient Primal-Dual Methods for Constrained MDPs.” *Journal of Machine Learning Research*, 26(256), 1–76, 2025.
Or
Eitan Altman. *Constrained Markov Decision Processes*. Chapman & Hall/CRC, 1999.
3. Wei Xiao, Christos G. Cassandras, and Calin Belta. *Safe Autonomy with Control Barrier Functions: Theory and Applications*. Springer, 2023.
4. Aaron D. Ames, Xiangru Xu, Jessie W. Grizzle, and Paulo Tabuada. “Control Barrier Function Based Quadratic Programs for Safety Critical Systems.” *IEEE Transactions on Automatic Control*, 62(8), 3861–3876, 2017.
5. Somil Bansal, Mo Chen, Sylvia Herbert, and Claire J. Tomlin. “Hamilton-jacobi reachability: A brief overview and recent advances.” *IEEE 56th annual conference on decision and control (CDC)*, 2242–2253, 2017.
6. Anastasios N. Angelopoulos and Stephen Bates. “Conformal Prediction: A Gentle Introduction.” *Foundations and Trends in Machine Learning*, 16(4), 494–591, 2023.
7. Anastasios N. Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, and Tal Schuster. “Conformal Risk Control.” *International Conference on Learning Representations*, 2024.

9. Any other Remarks

Dated: 30 Sept 2026

Proposer: Dr. Ashwin Verma

Dated:

DPGC Convener:

The course is approved / not approved

Chairman SUGC/SPGC:

Dated: