The Hong Kong University of Science and Technology

UG Course Syllabus Template

Stochastic Models
IEDA 3250
3 Credits
Prerequisites: (IEDA 2520 and IEDA 2540) or (MATH 2411 and MATH 2421)

Name: Wei You

Email: weiyou@ust.hk

Office Hours: Weekly office hours will be announced on Canvas; appointments in Room 5559D.

Course Description

Many operational systems evolve over time under uncertainty: customers arrive and wait, machines fail and are repaired, inventories rise and fall, and decisions affect future states. This course develops a unified modeling toolkit for such systems. It covers probability review, discrete-time Markov chains, Poisson processes, continuous-time Markov chains, birth-death processes, queueing models, reliability and inventory applications, and introductory Markov decision processes. Students learn to translate a system description into a stochastic model, derive transient and steady-state performance measures, interpret operational trade-offs, and validate analyses using Python/Jupyter notebooks. Instruction combines lectures, worked examples, tutorials, problem solving, and computational modeling.

Intended Learning Outcomes (ILOs)

By the end of this course, students should be able to:

1. Explain the assumptions and core concepts underlying Poisson processes, discrete- and continuous-time Markov chains, queueing systems, reliability and inventory models, and Markov decision processes.
2. Formulate real operational systems as appropriate stochastic models by defining states, transitions or rates, rewards or costs, and performance measures.
3. Analyze transient and steady-state behavior using transition matrices, first-step equations, balance equations, and Bellman recursions.
4. Evaluate service, reliability, inventory, and capacity trade-offs and communicate conclusions together with the assumptions and limitations of the model.
5. Implement and validate numerical or simulation-based analyses using Python and present reproducible results.

Week	Topic	Key concepts and applications
1	Introduction to modeling under uncertainty	Illustrative examples and basic concepts of stochastic models.
2	Discrete-time Markov chains: foundations	State spaces, transition matrices, sample paths, n-step probabilities, and Chapman-Kolmogorov equations.
3	Discrete-time Markov chains: long-run behavior	Communicating classes, recurrence and transience, periodicity, stationary distributions, and long-run averages.
4	First-step analysis and applications	Hitting probabilities and expected times; absorbing chains; introductory reliability and inventory examples.
5	Poisson processes	Exponential interarrival times, count distributions, thinning, superposition, and conditional arrival times.

Week	Topic	Key concepts and applications
6	Continuous-time Markov chains	Rate diagrams, generators, holding times, embedded chains, and Kolmogorov equations.
7	Birth-death processes and steady state	Global and detailed balance; stationary distributions; modeling applications.
8	Queueing fundamentals	Kendall notation, Little's law, PASTA, and the M/M/1 queue.
9	Queueing design and capacity	Finite-capacity and multi-server queues; Erlang loss and delay models; service-capacity trade-offs.
10	Reliability and inventory models	Repairable systems, availability, simple stock-control models, and a computational laboratory.
11	Markov decision processes I	States, actions, transition probabilities, rewards and costs, finite-horizon Bellman recursion.
12	Markov decision processes II	Discounted and average reward criteria; policy evaluation and improvement; operational applications.
13	Integrated modeling cases and review	Selecting an appropriate model, validating assumptions, interpreting results, and cumulative review.

Assessment and Grading

This course will be assessed using criterion-referencing and grades will not be assigned using a curve. Detailed rubrics are provided below. All assessed work is individual unless explicitly stated otherwise.

Assessments:

Assessment Task	Contribution to Overall Course Grade (%)	Due Date
Homework	20%	Bi-weekly
Quiz 1	25%	29/09/2026
Quiz 2	25%	03/11/2026
Final examination	30%	University examination period

** Marks and feedback for assessed tasks completed during the teaching term will normally be released within two weeks. Final examination results follow University procedures.*

- **Homework (20%):** Bi-weekly individual problem sets assess model recognition, formulation, analytical derivation, numerical calculation, and interpretation. Selected assignments may include short Python components for numerical or simulation-based validation. Task-specific instructions and rubrics will be posted on Canvas.
- **Quiz 1 (25%) and Quiz 2 (25%):** Closed-book, in-person assessments emphasizing model formulation, analytical derivation, and interpretation of the material covered up to each quiz date. One double-sided A4 formula sheet and a non-programmable calculator are permitted unless otherwise announced.
- **Final examination (30%):** A cumulative examination assessing integrated model selection, analysis, and interpretation. Permitted materials and examination arrangements will be confirmed on Canvas in advance.

Mapping of Course ILOs to Assessment Tasks

Assessed Task	Mapped ILOs	Explanation
Homework	ILO1, ILO2, ILO3, ILO4, ILO5	Bi-weekly problem sets assess understanding of stochastic-model concepts, model formulation, analytical solution methods, interpretation of operational trade-offs and model limitations, and reproducible numerical or simulation analysis.
Quiz 1	ILO1, ILO2, ILO3	Assesses probability foundations, discrete-time Markov-chain formulation and analysis, and Poisson-process calculations.
Quiz 2	ILO1, ILO2, ILO3, ILO4	Assesses concepts and assumptions, model formulation, transient and stead interpretation for the middle part of the course.

Assessed Task	Mapped ILOs	Explanation
Final examination	ILO1, ILO2, ILO3, ILO4	Assesses cumulative understanding, model selection and formulation, integrated analysis, and interpretation of stochastic systems in familiar and moderately unfamiliar settings.

Grading Rubrics

The rubrics below explain how marks are assigned within each assessed task. They are applied criterion by criterion and are distinct from the final grade descriptors, which summarize overall course performance. Exact question-level marks and task-specific instructions printed on an assessment take precedence over the default allocations below.

Rubric A: Analytical Questions in Homework, Quizzes, and the Final Examination

The percentages indicate a typical allocation across a multi-part analytical task. A question may assess only a subset of the criteria; its printed mark allocation governs.

Criterion	Excellent / Full or Near-Full Credit	Good / Minor Deductions	Satisfactory / Partial Credit	Marginal / Limited Credit	Fail / No Credible Evidence
Conceptual understanding and assumptions *Typical 15%; ILO1*	Identifies all relevant stochastic concepts and states material assumptions precisely; explains their implications and domain of validity.	Concepts and assumptions are essentially correct; one minor omission or imprecision does not affect the solution.	Recognizes the main concept, but assumptions are incomplete or weakly explained; a minor misconception affects part of the work.	Shows limited understanding; several key assumptions are missing or misapplied, so only part of the work is relevant.	Uses no relevant concept or fundamentally misunderstands the stochastic setting.
Model formulation *Typical 20%; ILO2*	Defines states, parameters, transitions or rates, initial or boundary conditions, rewards or costs, and the target measure unambiguously; the model fits the system.	Chooses an appropriate model with a minor missing definition or notation issue that is readily corrected.	The core model is recognizable, but one important component is incomplete or ambiguous; it supports only part of the requested analysis.	Major model elements are missing or inconsistent; the formulation is weakly connected to the stated system.	Provides no coherent model, or the model is incompatible with the problem.
Method selection and derivation *Typical 35%; ILO3*	Selects an appropriate method, develops a logically complete derivation, justifies nontrivial steps, and handles relevant cases or boundaries.	Uses the correct method and a largely complete derivation; only a minor gap or local conceptual error remains.	Uses a plausible method and makes substantial progress, but the derivation is incomplete or contains one material, recoverable error.	Writes some relevant equations but does not carry the method through; major logical gaps remain.	Uses no valid method or provides irrelevant work.
Accuracy and completeness *Typical 15%; ILO3*	Calculations are correct and complete; final results state the required conditions or units and are checked against probabilities, boundaries, or limiting cases.	The main result is correct; minor algebraic, numerical, or notation errors do not change the conclusion.	Produces a partly correct result, but several local errors or omitted cases reduce reliability.	Major computational errors make the answer unreliable despite some correct intermediate steps.	Provides no credible calculation or result.
Interpretation and communication *Typical 15%; ILO4 when assessed*	Answers the operational question directly, explains the result in system terms, and discusses relevant trade-offs, assumptions, or limitations using clear notation.	Interpretation is correct and presentation is clear; discussion of trade-offs or limitations is brief.	Provides a basic interpretation, but the link to the system, assumptions, or requested conclusion is incomplete or imprecise.	Mostly reports formulas or numbers; reasoning is difficult to follow and has little connection to the operational question.	Provides no relevant interpretation or an unsupported or unreadable answer.

Rubric B: Computational Components of Homework (when assigned)

This rubric applies when Python/Jupyter, numerical computation, or simulation is explicitly required. Analytical derivations accompanying the code are also assessed under Rubric A where relevant.

Rubric levels are not direct letter grades for the task or course. Work between two descriptors receives intermediate marks. A correct final answer without the required model or derivation does not receive full credit. Consistent follow-through after an earlier error may receive appropriate method credit when subsequent steps are logically valid. Blank, irrelevant, or non-reproducible work receives zero for the affected criterion.

Criterion	Excellent / Full or Near-Full Credit	Good / Minor Deductions	Satisfactory / Partial Credit	Marginal / Limited Credit	Fail / No Credible Evidence
Problem and experiment design *15%; ILO2, ILO5*	Defines the stochastic model, computational objective, inputs, outputs, parameter ranges, and performance metrics; the design directly addresses the question.	The design is sound and relevant; one minor input, metric, or justification is omitted.	The basic design is workable, but parameters, metrics, or rationale are only partly specified.	The design is poorly specified and cannot fully answer the question.	Provides no coherent computational design.
Analytical or numerical method *20%; ILO3, ILO5*	Implements the correct formulas or algorithms and chooses appropriate horizon, sample size, initialization, and numerical procedure; material choices are justified.	The core method is correct; a minor technical or efficiency issue does not affect the main conclusion.	The main method is recognizable, but technical choices are weakly justified or errors affect some results.	Major methodological errors or inappropriate settings make the results unreliable.	Implements no valid method.
Implementation and reproducibility *25%; ILO5*	The notebook runs from start to finish in a clean environment; dependencies, seeds, and data are specified; code is readable and all reported outputs are reproducible.	The notebook reproduces the main results with only minor organization, documentation, or efficiency weaknesses.	The notebook runs after small manual fixes or undocumented steps; code is understandable but only partly reproducible.	Substantial intervention, missing files or cells, or unclear code prevents reliable reproduction.	Code is absent, nonfunctional, or unrelated to the submitted results.
Validation and diagnostics *25%; ILO4, ILO5*	Checks results against theory, benchmarks, or special cases; reports convergence or uncertainty; investigates discrepancies and sensitivity.	Includes appropriate checks or comparisons; diagnostic or sensitivity analysis is limited.	Provides one basic check, but discrepancies or uncertainty are only weakly examined.	Validation is superficial or inappropriate; implausible outputs are not addressed.	Provides no validation, or conclusions contradict evident diagnostics.
Interpretation and presentation *15%; ILO4, ILO5*	Uses concise, well-labeled tables or figures and a coherent narrative; conclusions follow from the evidence and connect to operational decisions and limitations.	Results and conclusions are clear and correct; minor labeling or interpretive weaknesses remain.	The main results are visible and broadly interpreted, but labels, comparisons, or explanation are incomplete.	Outputs are cluttered or unclear; conclusions are weakly supported or disconnected from the question.	Provides no usable presentation or relevant interpretation.

Final Grade Descriptors:

Grade	Short Description	Elaboration on Subject Grading Description
A	Excellent Performance	Demonstrates a comprehensive and connected understanding of the course. Correctly formulates and analyzes standard and unfamiliar stochastic systems, selects appropriate methods, performs accurate calculations, validates results, and provides insightful operational interpretation with clear communication.
B	Good Performance	Shows strong knowledge of the main models and competent problem-solving. Correctly handles standard problems and many moderately unfamiliar settings, with only minor technical or interpretive weaknesses.
C	Satisfactory Performance	Possesses adequate knowledge of core concepts and can solve familiar problems using appropriate methods. Analysis and computation are generally serviceable but may show gaps in integration, validation, or interpretation.
D	Marginal Pass	Meets the minimum threshold for the course. Shows partial understanding of basic models and can complete some routine calculations, but performance is inconsistent and contains significant conceptual, computational, or communication weaknesses.
F	Fail	Demonstrates insufficient understanding of the fundamental models and methods. Cannot reliably formulate or analyze basic stochastic systems, and submitted work does not meet the minimum standards of accuracy, reasoning, or interpretation.

Course AI Policy

Generative AI tools may be used for ungraded study, concept explanations, brainstorming, code debugging, and language editing. For graded assignments, AI use is permitted only with a brief disclosure naming the tool and describing how it was used. Students remain responsible for verifying every derivation, model, result, citation, and line of submitted code, and must be able to explain their work. Generative AI may not be used to produce answers for quizzes, or the final examination. No electronic devices are permitted during in-person assessments except approved calculators.

Communication and Feedback

Canvas is the authoritative channel for announcements, materials, deadlines, and assessment instructions. Mastery checks provide immediate or prompt feedback. Quiz and assignment feedback will normally include a rubric or annotated solution, and common errors will be reviewed in lectures or tutorials. Students who have questions about feedback or marks should consult the instructor within five working days after the feedback is received.

Resubmission Policy

Routine resubmission or reassessment for additional credit is not permitted. Requests concerning a possible marking or administrative error should be made within five working days after feedback is released.

Required Texts and Materials

Primary text:

Sheldon M. Ross, Introduction to Probability Models, 13th edition, Academic Press, 2023.

Course materials: Lecture notes, worked examples, tutorial problems, and Python/Jupyter notebooks distributed through Canvas.

Academic Integrity

Students are expected to adhere to the university's academic integrity policy. Students are expected to uphold HKUST's Academic Honor Code and to maintain the highest standards of academic integrity. The University has zero tolerance of academic misconduct. Please refer to Academic Integrity | HKUST – Academic Registry for the University's definition of plagiarism and ways to avoid cheating and plagiarism.