

reinforcement learning and stochastic optimization

Reinforcement Learning and Stochastic Optimization: Unlocking Intelligent Decision-Making

reinforcement learning and stochastic optimization are two powerful concepts that have gained significant traction in the fields of artificial intelligence and machine learning. Both approaches center around making decisions in uncertain environments, but they tackle the problem from slightly different angles. When combined, these techniques offer a robust framework for solving complex optimization problems where uncertainty and dynamic interaction play crucial roles. In this article, we'll explore what reinforcement learning and stochastic optimization entail, how they relate, and why their synergy is vital for advancing intelligent systems.

Understanding Reinforcement Learning

Reinforcement learning (RL) is a type of machine learning where an agent learns how to make decisions by interacting with its environment. Unlike supervised learning, where the model learns from labeled data, RL relies on feedback in the form of rewards or penalties. The agent's goal is to maximize the cumulative reward over time by choosing actions that lead to the best outcomes.

The Core Components of Reinforcement Learning

At the heart of reinforcement learning lies the Markov Decision Process (MDP), which formalizes the environment in terms of states, actions, rewards, and transitions. Here's a quick breakdown:

- **States (S):** The different situations or configurations the agent can be in.
- **Actions (A):** The possible decisions or moves the agent can make.
- **Rewards (R):** Feedback signals that indicate the quality of an action in a given state.
- **Transition Dynamics:** The probabilities of moving from one state to another after taking a specific action.

The agent uses a policy—a strategy mapping states to actions—to decide what to do next. Over time, through trial and error, it learns an optimal policy that maximizes expected rewards, often employing techniques like Q-learning or policy gradients.

Applications Where Reinforcement Learning Shines

Reinforcement learning is particularly suited for problems where sequential decisions matter, and the environment is either unknown or changing. Some real-world applications include:

- **Robotics:** Teaching robots to navigate or manipulate objects.
- **Game Playing:** Algorithms that master games such as chess, Go, or video games.
- **Recommendation Systems:** Personalizing content dynamically based on user interactions.
- **Autonomous Vehicles:** Making driving decisions in unpredictable road conditions.

The Essence of Stochastic Optimization

While reinforcement learning focuses on learning optimal policies, stochastic optimization is a broader mathematical framework aimed at optimizing functions that involve randomness or uncertainty. It's often employed when the objective function is noisy, non-deterministic, or too complex for traditional deterministic optimization methods.

What Makes an Optimization Problem “Stochastic”?

In classical optimization, you might try to find the minimum or maximum of a well-defined function. However, in many real-world scenarios, the function to be optimized isn't fixed or known precisely. Instead, it may depend on random variables or probabilistic events, making the problem stochastic by nature.

For example, consider the task of optimizing the supply chain in a company where customer demand is uncertain. The objective function—say, total cost—depends on this stochastic demand, so the optimization process must account for variability and unpredictability.

Common Techniques in Stochastic Optimization

To tackle such problems, various algorithms have been developed, including:

- **Stochastic Gradient Descent (SGD):** Used extensively in training neural networks, it updates parameters based on noisy gradient estimates from random samples.

- **Simulated Annealing:** A probabilistic technique mimicking the cooling of metals to escape local minima.
- **Genetic Algorithms:** Inspired by natural selection, using randomness to explore the solution space.
- **Stochastic Approximation:** Methods that iteratively refine estimates using noisy observations.

These methods balance exploration and exploitation to find good solutions in uncertain environments.

How Reinforcement Learning and Stochastic Optimization Intersect

Reinforcement learning and stochastic optimization are deeply intertwined. At its core, reinforcement learning often involves solving stochastic optimization problems, especially when the environment is uncertain or the reward signals are noisy.

Reinforcement Learning as a Stochastic Optimization Problem

In RL, the objective is to find a policy that maximizes expected cumulative reward. Because the environment's dynamics and rewards are usually uncertain or partially observable, this optimization is inherently stochastic. The agent must optimize over a distribution of possible outcomes, which requires stochastic optimization methods.

For example, policy gradient algorithms optimize the parameters of a policy by estimating gradients of the expected reward. These estimates are noisy because they rely on sampled trajectories, making the process a stochastic optimization.

Using Stochastic Optimization to Enhance RL Algorithms

Many reinforcement learning algorithms borrow from stochastic optimization to improve learning efficiency and stability:

- SGD and its variants (Adam, RMSProp) are standard optimizers for training neural networks within RL agents.
- Stochastic methods help in exploration strategies, enabling agents to avoid local

optima by injecting randomness.

- Batching and experience replay use stochastic sampling to break correlations and reduce variance in updates.

This synergy allows reinforcement learning agents to learn effectively even in high-dimensional or noisy environments.

Challenges and Practical Considerations

Despite their strengths, combining reinforcement learning and stochastic optimization presents unique challenges.

Sample Efficiency and Exploration

RL agents often require a vast number of interactions with the environment to learn good policies. This can be computationally expensive or impractical in real-world settings. Stochastic optimization helps by improving gradient estimates and reducing variance, but balancing exploration (trying new actions) and exploitation (using known good actions) remains tricky.

Dealing with Non-Stationarity

Both RL environments and stochastic optimization problems can be non-stationary, meaning their properties change over time. For instance, user preferences in a recommendation system evolve, or the dynamics of a physical system may shift. Algorithms must adapt to these changes, which often requires online learning and continual optimization strategies.

Stability and Convergence

Stochastic optimization methods can suffer from instability due to noisy gradients or poor hyperparameter choices. In reinforcement learning, instability can lead to divergent policies or suboptimal performance. Techniques such as target networks, entropy regularization, and adaptive learning rates help mitigate these issues.

Emerging Trends and Future Directions

The intersection of reinforcement learning and stochastic optimization continues to be a

fertile area for research, with exciting developments on the horizon.

Deep Reinforcement Learning and Stochastic Optimization

Deep learning has revolutionized reinforcement learning by enabling agents to handle high-dimensional inputs like images or raw sensor data. Training deep RL agents relies heavily on stochastic optimization algorithms, pushing research toward more efficient and robust optimizers that can handle the scale and complexity.

Meta-Learning and Automated Optimization

Meta-learning seeks to create agents that learn how to learn, adapting quickly to new tasks. This involves stochastic optimization at multiple levels — optimizing not only policies but also the learning algorithms themselves. Reinforcement learning agents empowered by meta-optimization strategies promise faster adaptation and better generalization.

Applications in Finance, Healthcare, and Beyond

The combination of reinforcement learning and stochastic optimization is increasingly applied in domains like:

- Algorithmic trading to optimize portfolios under uncertainty.
- Personalized medicine, where treatment decisions adapt to patient responses.
- Energy management systems optimizing power grids amid fluctuating demand and supply.

These real-world problems perfectly showcase the strengths of integrating decision-making under uncertainty with sophisticated optimization tools.

Exploring the intersection of reinforcement learning and stochastic optimization reveals a vibrant landscape of methods designed to make intelligent decisions in complex, uncertain environments. As these fields evolve, they continue to push the boundaries of what machines can achieve, transforming industries and opening pathways to smarter, more adaptive technologies.

Frequently Asked Questions

What is the relationship between reinforcement learning and stochastic optimization?

Reinforcement learning (RL) is a framework for learning optimal decision-making policies through interaction with an environment, often modeled as a Markov decision process. Stochastic optimization refers to techniques for optimizing objective functions that involve randomness or uncertainty. RL problems are commonly solved using stochastic optimization methods because the environment's dynamics and rewards are often uncertain and noisy.

How do policy gradient methods in reinforcement learning utilize stochastic optimization?

Policy gradient methods optimize the parameters of a policy directly by estimating gradients of expected rewards with respect to those parameters. Since the reward signals are stochastic due to environment randomness and policy exploration, these methods use stochastic gradient ascent techniques to iteratively improve the policy.

What are some common stochastic optimization algorithms used in reinforcement learning?

Common stochastic optimization algorithms in reinforcement learning include stochastic gradient descent (SGD), Adam optimizer, natural gradient methods, and evolutionary strategies. These algorithms help update policy or value function parameters in the presence of noisy gradient estimates derived from sampled interactions.

How does stochastic optimization help address exploration-exploitation trade-offs in reinforcement learning?

Stochastic optimization introduces randomness in policy updates or parameter sampling, which can encourage exploration of the environment. By balancing the variance in updates and convergence speed, stochastic optimization methods help reinforcement learning agents explore new strategies while still exploiting known rewarding behaviors.

Can reinforcement learning be applied to stochastic optimization problems outside typical control tasks?

Yes, reinforcement learning can be applied to solve general stochastic optimization problems where the objective or constraints have uncertainty. For example, RL has been used for portfolio optimization, supply chain management, and adaptive resource allocation, where the environment's stochastic nature requires learning robust decision-making policies.

What challenges arise when combining reinforcement learning with stochastic optimization techniques?

Challenges include high variance in gradient estimates leading to unstable learning, sample inefficiency due to noisy feedback, difficulty in tuning hyperparameters, and convergence issues in non-stationary or high-dimensional environments. Addressing these challenges often requires variance reduction techniques, better exploration strategies, and more efficient data utilization.

Additional Resources

Reinforcement Learning and Stochastic Optimization: Exploring the Intersection of Adaptive Decision-Making and Uncertainty Management

reinforcement learning and stochastic optimization have emerged as pivotal methodologies in the evolving landscape of artificial intelligence and operations research. Each discipline addresses complex decision-making problems, often under uncertainty, yet their intersection offers profound capabilities for developing adaptive systems that learn optimal strategies through interaction with dynamic environments. Understanding how reinforcement learning (RL) leverages stochastic optimization principles not only advances theoretical insights but also enhances practical applications ranging from robotics to finance.

Understanding Reinforcement Learning and Stochastic Optimization

Reinforcement learning is a branch of machine learning focused on how agents ought to take actions in an environment to maximize cumulative rewards. Unlike supervised learning, RL does not rely on labeled input-output pairs but instead learns through trial-and-error interactions, receiving feedback in the form of rewards or penalties. This feedback loop enables an agent to develop policies that map states to actions, optimizing long-term benefits.

Stochastic optimization, on the other hand, involves optimizing objective functions that are inherently noisy or uncertain. These problems arise when the function evaluations or constraints depend on random variables or are subject to measurement errors. Stochastic optimization algorithms aim to find solutions that perform well on average or with high probability, rather than optimizing a deterministic objective.

The Confluence of Learning and Optimization Under Uncertainty

Reinforcement learning intrinsically involves stochastic optimization because the environment's dynamics and the reward signals are often probabilistic. For example, an

agent navigating a maze may face uncertain outcomes from its actions due to sensor noise or unpredictable changes in the environment. Consequently, RL algorithms frequently employ stochastic optimization techniques to update policies or value functions.

Consider policy gradient methods, which optimize parameterized policies by estimating gradients of expected rewards with respect to policy parameters. These gradients are inherently noisy due to sampling from stochastic environments, necessitating stochastic optimization approaches such as stochastic gradient descent (SGD) or its variants to converge to optimal policies.

Key Algorithms and Methodologies at the Intersection

Several reinforcement learning algorithms explicitly integrate stochastic optimization principles to handle uncertainty and improve learning efficiency.

Value-Based Methods and Stochastic Approximation

Algorithms like Q-learning and SARSA estimate value functions representing expected future rewards. These methods update value estimates incrementally using noisy samples from the environment, an approach known as stochastic approximation. The convergence properties of these algorithms depend heavily on the stochastic optimization framework, ensuring that estimated values approach the true values despite randomness in observations.

Policy Gradient and Actor-Critic Approaches

Policy gradient methods optimize policies directly by ascending the gradient of expected returns. Because the gradient estimates are stochastic, these methods rely on stochastic optimization theory to manage variance and maintain stable learning. Actor-critic algorithms combine value function estimation (critic) with policy optimization (actor), using stochastic optimization to update both components simultaneously, balancing bias and variance effectively.

Evolutionary Strategies and Derivative-Free Optimization

Some reinforcement learning scenarios employ evolutionary strategies, which are stochastic optimization techniques inspired by natural selection. These methods do not require gradient information and instead evaluate populations of policy candidates, selecting and mutating them based on performance. This approach is particularly useful in high-dimensional or non-differentiable environments where traditional gradient-based

methods struggle.

Applications Leveraging the Synergy Between Reinforcement Learning and Stochastic Optimization

The fusion of reinforcement learning and stochastic optimization has catalyzed breakthroughs across diverse domains.

Robotics and Autonomous Systems

Robotic control requires adapting to uncertain and dynamic conditions. Reinforcement learning algorithms, supported by stochastic optimization, enable robots to learn dexterous manipulation or navigation policies from trial-and-error experiences. The stochastic nature of sensor data and actuation errors makes robust optimization critical to achieving reliable performance.

Finance and Portfolio Management

Financial markets are paradigmatic stochastic environments, with asset prices fluctuating unpredictably. Reinforcement learning models, using stochastic optimization to handle noisy reward signals and market volatility, have been applied to portfolio optimization, algorithmic trading, and risk management, seeking strategies that maximize returns while controlling downside risks.

Supply Chain and Operations Management

In complex supply chains, decision-making under uncertainty is fundamental. Reinforcement learning integrated with stochastic optimization techniques assists in inventory management, demand forecasting, and dynamic pricing, enabling systems to adaptively optimize operations amid fluctuating supply, demand, and costs.

Challenges and Considerations in Combining Reinforcement Learning with Stochastic Optimization

While their integration is powerful, several challenges persist.

- **Sample Efficiency:** Reinforcement learning often requires large amounts of interaction data, which can be costly or impractical in real-world systems. Stochastic optimization methods must be designed to maximize learning from limited samples.
- **Exploration vs. Exploitation:** Balancing the exploration of unknown strategies and exploitation of learned policies remains a fundamental difficulty, magnified by stochastic reward structures.
- **Convergence Guarantees:** The stochastic nature of both learning and optimization introduces noise and instability, complicating theoretical guarantees on convergence and performance.
- **Computational Complexity:** High-dimensional state and action spaces increase computational demands, requiring scalable stochastic optimization algorithms to maintain practical feasibility.

Advancements in Optimization Techniques Benefiting Reinforcement Learning

Recent developments in stochastic optimization have directly influenced reinforcement learning capabilities:

1. **Variance Reduction Methods:** Techniques such as control variates, importance sampling, and actor-critic architectures reduce the variance in gradient estimates, leading to faster and more stable learning.
2. **Adaptive Learning Rates:** Algorithms like Adam or RMSProp dynamically adjust learning rates during optimization, improving convergence speed and robustness.
3. **Trust Region Methods:** By constraining policy updates within a trust region, methods like TRPO (Trust Region Policy Optimization) maintain policy stability despite stochastic updates.
4. **Bayesian Optimization:** Used for hyperparameter tuning, Bayesian methods guide stochastic searches efficiently over uncertain objective landscapes.

The Future Trajectory of Reinforcement Learning and Stochastic Optimization

As artificial intelligence systems increasingly operate in real-world, uncertain environments, the interplay between reinforcement learning and stochastic optimization will deepen. Emerging trends include the integration of model-based RL with stochastic optimization to

improve sample efficiency, the use of meta-learning to adapt optimization strategies dynamically, and leveraging advances in hardware acceleration to manage computational complexity.

Furthermore, the exploration of safe reinforcement learning, where stochastic optimization techniques help ensure adherence to safety constraints, marks a crucial frontier for deploying autonomous systems in sensitive settings such as healthcare and transportation.

In summary, reinforcement learning and stochastic optimization form a complementary pair addressing the dual challenges of learning from interaction and managing uncertainty. Their continuous evolution promises to unlock increasingly sophisticated, adaptive, and resilient decision-making systems across a spectrum of industries and applications.

Reinforcement Learning And Stochastic Optimization

reinforcement learning and stochastic optimization: *Reinforcement Learning and Stochastic Optimization* Warren B. Powell, 2022-04-25 REINFORCEMENT LEARNING AND STOCHASTIC OPTIMIZATION Clearing the jungle of stochastic optimization Sequential decision problems, which consist of “decision, information, decision, information,” are ubiquitous, spanning virtually every human activity ranging from business applications, health (personal and public health, and medical decision making), energy, the sciences, all fields of engineering, finance, and e-commerce. The diversity of applications attracted the attention of at least 15 distinct fields of research, using eight distinct notational systems which produced a vast array of analytical tools. A byproduct is that powerful tools developed in one community may be unknown to other communities. Reinforcement Learning and Stochastic Optimization offers a single canonical framework that can model any sequential decision problem using five core components: state variables, decision variables, exogenous information variables, transition function, and objective function. This book highlights twelve types of uncertainty that might enter any model and pulls together the diverse set of methods for making decisions, known as policies, into four fundamental classes that span every method suggested in the academic literature or used in practice. Reinforcement Learning and Stochastic Optimization is the first book to provide a balanced treatment of the different methods for modeling and solving sequential decision problems, following the style used by most books on machine learning, optimization, and simulation. The presentation is designed for readers with a course in probability and statistics, and an interest in modeling and applications. Linear programming is occasionally used for specific problem classes. The book is designed for readers who are new to the field, as well as those with some background in optimization under uncertainty. Throughout this book, readers will find references to over 100 different applications, spanning pure learning problems, dynamic resource allocation problems, general state-dependent problems, and hybrid learning/resource allocation problems such as those that arose in the COVID pandemic. There are 370 exercises, organized into seven groups, ranging from review questions, modeling, computation, problem solving, theory, programming exercises and a diary problem that a reader chooses at the beginning of the book, and which is used as a basis for questions throughout the rest of the book.

reinforcement learning and stochastic optimization: *Stochastic Optimization Methods for Policy Evaluation in Reinforcement Learning* Yi Zhou, Shaocong Ma, 2024-07-11 This monograph introduces various value-based approaches for solving the policy evaluation problem in the online reinforcement learning (RL) scenario, which aims to learn the value function associated with a specific policy under a single Markov decision process (MDP). Approaches vary depending on

whether they are implemented in an on-policy or off-policy manner. In on-policy settings, where the evaluation of the policy is conducted using data generated from the same policy that is being assessed, classical techniques such as TD(0), TD(λ), and their extensions with function approximation or variance reduction are employed in this setting. For off-policy evaluation, where samples are collected under a different behavior policy, this monograph introduces gradient-based two-timescale algorithms like GTD2, TDC, and variance-reduced TDC. These algorithms are designed to minimize the mean-squared projected Bellman error (MSPBE) as the objective function. This monograph also discusses their finite-sample convergence upper bounds and sample complexity.

reinforcement learning and stochastic optimization: Reinforcement Learning for Discrete Stochastic Optimization Deniz Daniel Preil, 2022

reinforcement learning and stochastic optimization: Reinforcement Learning and Optimal Control Dimitri Bertsekas, 2019-07-01 This book considers large and challenging multistage decision problems, which can be solved in principle by dynamic programming (DP), but their exact solution is computationally intractable. We discuss solution methods that rely on approximations to produce suboptimal policies with adequate performance. These methods are collectively known by several essentially equivalent names: reinforcement learning, approximate dynamic programming, neuro-dynamic programming. They have been at the forefront of research for the last 25 years, and they underlie, among others, the recent impressive successes of self-learning in the context of games such as chess and Go. Our subject has benefited greatly from the interplay of ideas from optimal control and from artificial intelligence, as it relates to reinforcement learning and simulation-based neural network methods. One of the aims of the book is to explore the common boundary between these two fields and to form a bridge that is accessible by workers with background in either field. Another aim is to organize coherently the broad mosaic of methods that have proved successful in practice while having a solid theoretical and/or logical foundation. This may help researchers and practitioners to find their way through the maze of competing ideas that constitute the current state of the art. This book relates to several of our other books: Neuro-Dynamic Programming (Athena Scientific, 1996), Dynamic Programming and Optimal Control (4th edition, Athena Scientific, 2017), Abstract Dynamic Programming (2nd edition, Athena Scientific, 2018), and Nonlinear Programming (Athena Scientific, 2016). However, the mathematical style of this book is somewhat different. While we provide a rigorous, albeit short, mathematical account of the theory of finite and infinite horizon dynamic programming, and some fundamental approximation methods, we rely more on intuitive explanations and less on proof-based insights. Moreover, our mathematical requirements are quite modest: calculus, a minimal use of matrix-vector algebra, and elementary probability (mathematically complicated arguments involving laws of large numbers and stochastic convergence are bypassed in favor of intuitive explanations). The book illustrates the methodology with many examples and illustrations, and uses a gradual expository approach, which proceeds along four directions: (a) From exact DP to approximate DP: We first discuss exact DP algorithms, explain why they may be difficult to implement, and then use them as the basis for approximations. (b) From finite horizon to infinite horizon problems: We first discuss finite horizon exact and approximate DP methodologies, which are intuitive and mathematically simple, and then progress to infinite horizon problems. (c) From deterministic to stochastic models: We often discuss separately deterministic and stochastic problems, since deterministic problems are simpler and offer special advantages for some of our methods. (d) From model-based to model-free implementations: We first discuss model-based implementations, and then we identify schemes that can be appropriately modified to work with a simulator. The book is related and supplemented by the companion research monograph Rollout, Policy Iteration, and Distributed Reinforcement Learning (Athena Scientific, 2020), which focuses more closely on several topics related to rollout, approximate policy iteration, multiagent problems, discrete and Bayesian optimization, and distributed computation, which are either discussed in less detail or not covered at all in the present book. The author's website contains class notes, and a series of videolectures and

slides from a 2021 course at ASU, which address a selection of topics from both books.

reinforcement learning and stochastic optimization: Stochastic Optimization for Large-scale Machine Learning Vinod Kumar Chauhan, 2021-11-18 Advancements in the technology and availability of data sources have led to the 'Big Data' era. Working with large data offers the potential to uncover more fine-grained patterns and take timely and accurate decisions, but it also creates a lot of challenges such as slow training and scalability of machine learning models. One of the major challenges in machine learning is to develop efficient and scalable learning algorithms, i.e., optimization techniques to solve large scale learning problems. Stochastic Optimization for Large-scale Machine Learning identifies different areas of improvement and recent research directions to tackle the challenge. Developed optimisation techniques are also explored to improve machine learning algorithms based on data access and on first and second order optimisation methods. Key Features: Bridges machine learning and Optimisation. Bridges theory and practice in machine learning. Identifies key research areas and recent research directions to solve large-scale machine learning problems. Develops optimisation techniques to improve machine learning algorithms for big data problems. The book will be a valuable reference to practitioners and researchers as well as students in the field of machine learning.

reinforcement learning and stochastic optimization: Approximate Dynamic Programming Warren B. Powell, 2011-10-26 Praise for the First Edition Finally, a book devoted to dynamic programming and written using the language of operations research (OR)! This beautiful book fills a gap in the libraries of OR specialists and practitioners. —Computing Reviews This new edition showcases a focus on modeling and computation for complex classes of approximate dynamic programming problems Understanding approximate dynamic programming (ADP) is vital in order to develop practical and high-quality solutions to complex industrial problems, particularly when those problems involve making decisions in the presence of uncertainty. Approximate Dynamic Programming, Second Edition uniquely integrates four distinct disciplines—Markov decision processes, mathematical programming, simulation, and statistics—to demonstrate how to successfully approach, model, and solve a wide range of real-life problems using ADP. The book continues to bridge the gap between computer science, simulation, and operations research and now adopts the notation and vocabulary of reinforcement learning as well as stochastic search and simulation optimization. The author outlines the essential algorithms that serve as a starting point in the design of practical solutions for real problems. The three curses of dimensionality that impact complex problems are introduced and detailed coverage of implementation challenges is provided. The Second Edition also features: A new chapter describing four fundamental classes of policies for working with diverse stochastic optimization problems: myopic policies, look-ahead policies, policy function approximations, and policies based on value function approximations A new chapter on policy search that brings together stochastic search and simulation optimization concepts and introduces a new class of optimal learning strategies Updated coverage of the exploration exploitation problem in ADP, now including a recently developed method for doing active learning in the presence of a physical state, using the concept of the knowledge gradient A new sequence of chapters describing statistical methods for approximating value functions, estimating the value of a fixed policy, and value function approximation while searching for optimal policies The presented coverage of ADP emphasizes models and algorithms, focusing on related applications and computation while also discussing the theoretical side of the topic that explores proofs of convergence and rate of convergence. A related website features an ongoing discussion of the evolving fields of approximation dynamic programming and reinforcement learning, along with additional readings, software, and datasets. Requiring only a basic understanding of statistics and probability, Approximate Dynamic Programming, Second Edition is an excellent book for industrial engineering and operations research courses at the upper-undergraduate and graduate levels. It also serves as a valuable reference for researchers and professionals who utilize dynamic programming, stochastic programming, and control theory to solve problems in their everyday work.

reinforcement learning and stochastic optimization: First-order and Stochastic

Optimization Methods for Machine Learning Guanghui Lan, 2020-05-15 This book covers not only foundational materials but also the most recent progresses made during the past few years on the area of machine learning algorithms. In spite of the intensive research and development in this area, there does not exist a systematic treatment to introduce the fundamental concepts and recent progresses on machine learning algorithms, especially on those based on stochastic optimization methods, randomized algorithms, nonconvex optimization, distributed and online learning, and projection free methods. This book will benefit the broad audience in the area of machine learning, artificial intelligence and mathematical programming community by presenting these recent developments in a tutorial style, starting from the basic building blocks to the most carefully designed and complicated algorithms for machine learning.

reinforcement learning and stochastic optimization: *A Modern Approach to Teaching an Introduction to Optimization* Warren B. Powell, 2024-01-24 Designed for instructors looking to introduce the science of making good decisions. It should also be of interest to anyone who has already had a traditional course in optimization of any type.

reinforcement learning and stochastic optimization: *Handbook of Reinforcement Learning and Control* Kyriakos G. Vamvoudakis, Yan Wan, Frank L. Lewis, Derya Cansever, 2021-06-23 This handbook presents state-of-the-art research in reinforcement learning, focusing on its applications in the control and game theory of dynamic systems and future directions for related research and technology. The contributions gathered in this book deal with challenges faced when using learning and adaptation methods to solve academic and industrial problems, such as optimization in dynamic environments with single and multiple agents, convergence and performance analysis, and online implementation. They explore means by which these difficulties can be solved, and cover a wide range of related topics including: deep learning; artificial intelligence; applications of game theory; mixed modality learning; and multi-agent reinforcement learning. Practicing engineers and scholars in the field of machine learning, game theory, and autonomous control will find the Handbook of Reinforcement Learning and Control to be thought-provoking, instructive and informative.

reinforcement learning and stochastic optimization: *Adversarial Machine Learning* Aneesh Sreevallabh Chivukula, Xinghao Yang, Bo Liu, Wei Liu, Wanlei Zhou, 2023-03-06 A critical challenge in deep learning is the vulnerability of deep learning networks to security attacks from intelligent cyber adversaries. Even innocuous perturbations to the training data can be used to manipulate the behaviour of deep networks in unintended ways. In this book, we review the latest developments in adversarial attack technologies in computer vision; natural language processing; and cybersecurity with regard to multidimensional, textual and image data, sequence data, and temporal data. In turn, we assess the robustness properties of deep learning networks to produce a taxonomy of adversarial examples that characterises the security of learning systems using game theoretical adversarial deep learning algorithms. The state-of-the-art in adversarial perturbation-based privacy protection mechanisms is also reviewed. We propose new adversary types for game theoretical objectives in non-stationary computational learning environments. Proper quantification of the hypothesis set in the decision problems of our research leads to various functional problems, oracular problems, sampling tasks, and optimization problems. We also address the defence mechanisms currently available for deep learning models deployed in real-world environments. The learning theories used in these defence mechanisms concern data representations, feature manipulations, misclassifications costs, sensitivity landscapes, distributional robustness, and complexity classes of the adversarial deep learning algorithms and their applications. In closing, we propose future research directions in adversarial deep learning applications for resilient learning system design and review formalized learning assumptions concerning the attack surfaces and robustness characteristics of artificial intelligence applications so as to deconstruct the contemporary adversarial deep learning designs. Given its scope, the book will be of interest to Adversarial Machine Learning practitioners and Adversarial Artificial Intelligence researchers whose work involves the design and application of Adversarial Deep Learning.

reinforcement learning and stochastic optimization: *Comparative Analysis of Evolving*

Artificial Neural Network and Reinforcement Learning in Stochastic Optimization of Multi-reservoir Systems Amir Mohammad Moradi, 2017 Dynamic programming (DP) has been among the most popular techniques for solving multi-reservoir problems since the early 1960s. However, DP and DP-based methods suffer from two serious issues: namely, the curses of modelling and dimensionality. Later, reinforcement learning (RL) was introduced to overcome some deficiencies in the traditional DP mainly related to the curse of modelling, but it still encounters the curse of dimensionality in larger systems. Recently, the artificial neural network has emerged as an effective approach to solve stochastic optimization of reservoir systems with high flexibility. In this paper, we develop a single-step evolving artificial neural network (SENN) model that overcomes the curses of modelling and dimensionality in a multi-reservoir system. Furthermore, a novel efficient allocation technique is developed to ease the allocation of water among different users. A two-reservoir system in Karkheh Basin, Iran, is applied to derive and test the methods. Elitist-mutated particle swarm optimization is used to train the network. A comparison of the results with Q-learning shows the superiority of the SENN, especially during drought periods. Moreover, the SENN performs better in producing more hydropower energy in the system. Thus, the main contributions of this research are (1) development of SENN applications to multi-reservoir systems, (2) a comparative analysis between SENN and Q-learning especially in prolonged drought conditions and (3) a proposed efficient optimal allocation technique using the simulation method.

reinforcement learning and stochastic optimization: European Control Conference 1993, 1993-06-28 Proceedings of the European Control Conference 1993, Groningen, Netherlands, June 28 – July 1, 1993

reinforcement learning and stochastic optimization: Reinforcement Learning for Cyber-Physical Systems Chong Li, Meikang Qiu, 2019-02-22 Reinforcement Learning for Cyber-Physical Systems: with Cybersecurity Case Studies was inspired by recent developments in the fields of reinforcement learning (RL) and cyber-physical systems (CPSs). Rooted in behavioral psychology, RL is one of the primary strands of machine learning. Different from other machine learning algorithms, such as supervised learning and unsupervised learning, the key feature of RL is its unique learning paradigm, i.e., trial-and-error. Combined with the deep neural networks, deep RL become so powerful that many complicated systems can be automatically managed by AI agents at a superhuman level. On the other hand, CPSs are envisioned to revolutionize our society in the near future. Such examples include the emerging smart buildings, intelligent transportation, and electric grids. However, the conventional hand-programming controller in CPSs could neither handle the increasing complexity of the system, nor automatically adapt itself to new situations that it has never encountered before. The problem of how to apply the existing deep RL algorithms, or develop new RL algorithms to enable the real-time adaptive CPSs, remains open. This book aims to establish a linkage between the two domains by systematically introducing RL foundations and algorithms, each supported by one or a few state-of-the-art CPS examples to help readers understand the intuition and usefulness of RL techniques. Features Introduces reinforcement learning, including advanced topics in RL Applies reinforcement learning to cyber-physical systems and cybersecurity Contains state-of-the-art examples and exercises in each chapter Provides two cybersecurity case studies Reinforcement Learning for Cyber-Physical Systems with Cybersecurity Case Studies is an ideal text for graduate students or junior/senior undergraduates in the fields of science, engineering, computer science, or applied mathematics. It would also prove useful to researchers and engineers interested in cybersecurity, RL, and CPS. The only background knowledge required to appreciate the book is a basic knowledge of calculus and probability theory.

reinforcement learning and stochastic optimization: Decision Sciences Angel A. Juan, Javier Faulin, David Lopez-Lopez, 2025-01-30 This book constitutes the proceedings of the Second Decision Science Alliance International Summer Conference, DSA ISC 2024, held in Valencia, Spain, in June 2024. The 33 full papers and 38 short papers included in this book were carefully reviewed and selected from 101 submissions. At the core of DSA ISC'24 are in-depth discussions and analyses across a spectrum of technological domains. Notably, experts shared their knowledge on areas such

as Artificial Intelligence & Machine Learning, Mathematical Optimization, Operational Research & Management Science, Statistics, Simulation, and Decision Processes Analysis. Each of these areas represents a key aspect of decision science, contributing to the interdisciplinary nature of the conference.

reinforcement learning and stochastic optimization: *Reinforcement Learning Algorithms: Analysis and Applications* Boris Belousov, Hany Abdulsamad, Pascal Klink, Simone Parisi, Jan Peters, 2021-01-02 This book reviews research developments in diverse areas of reinforcement learning such as model-free actor-critic methods, model-based learning and control, information geometry of policy searches, reward design, and exploration in biology and the behavioral sciences. Special emphasis is placed on advanced ideas, algorithms, methods, and applications. The contributed papers gathered here grew out of a lecture course on reinforcement learning held by Prof. Jan Peters in the winter semester 2018/2019 at Technische Universität Darmstadt. The book is intended for reinforcement learning students and researchers with a firm grasp of linear algebra, statistics, and optimization. Nevertheless, all key concepts are introduced in each chapter, making the content self-contained and accessible to a broader audience.

reinforcement learning and stochastic optimization: Rollout, Policy Iteration, and Distributed Reinforcement Learning Dimitri Bertsekas, 2021-08-20 The purpose of this book is to develop in greater depth some of the methods from the author's Reinforcement Learning and Optimal Control recently published textbook (Athena Scientific, 2019). In particular, we present new research, relating to systems involving multiple agents, partitioned architectures, and distributed asynchronous computation. We pay special attention to the contexts of dynamic programming/policy iteration and control theory/model predictive control. We also discuss in some detail the application of the methodology to challenging discrete/combinatorial optimization problems, such as routing, scheduling, assignment, and mixed integer programming, including the use of neural network approximations within these contexts. The book focuses on the fundamental idea of policy iteration, i.e., start from some policy, and successively generate one or more improved policies. If just one improved policy is generated, this is called rollout, which, based on broad and consistent computational experience, appears to be one of the most versatile and reliable of all reinforcement learning methods. In this book, rollout algorithms are developed for both discrete deterministic and stochastic DP problems, and the development of distributed implementations in both multiagent and multiprocessor settings, aiming to take advantage of parallelism. Approximate policy iteration is more ambitious than rollout, but it is a strictly off-line method, and it is generally far more computationally intensive. This motivates the use of parallel and distributed computation. One of the purposes of the monograph is to discuss distributed (possibly asynchronous) methods that relate to rollout and policy iteration, both in the context of an exact and an approximate implementation involving neural networks or other approximation architectures. Much of the new research is inspired by the remarkable AlphaZero chess program, where policy iteration, value and policy networks, approximate lookahead minimization, and parallel computation all play an important role.

reinforcement learning and stochastic optimization: *Networks of Learning Automata* M.A.L. Thathachar, P.S. Sastry, 2003-10-31 Networks of Learning Automata: Techniques for Online Stochastic Optimization is a comprehensive account of learning automata models with emphasis on multiautomata systems. It considers synthesis of complex learning structures from simple building blocks and uses stochastic algorithms for refining probabilities of selecting actions. Mathematical analysis of the behavior of games and feedforward networks is provided. Algorithms considered here can be used for online optimization of systems based on noisy measurements of performance index. Also, algorithms that assure convergence to the global optimum are presented. Parallel operation of automata systems for improving speed of convergence is described. The authors also include extensive discussion of how learning automata solutions can be constructed in a variety of applications.

reinforcement learning and stochastic optimization: *Advances in Swarm Intelligence* Ying Tan, Yuhui Shi, Wenjian Luo, 2023-07-07 This two-volume set LNCS 13968 and 13969 constitutes

the proceedings of the 14th International Conference on Advances in Swarm Intelligence, ICSI 2023, which took place in Shenzhen, China, China, in July 2023. The theme of this year's conference was "Serving Life with Swarm Intelligence". The 81 full papers presented were carefully reviewed and selected from 170 submissions. The papers are organized into 12 cohesive sections covering major topics of swarm intelligence research and its development and applications. The papers of the first part cover topics such as: Swarm Intelligence Computing; Swarm Intelligence Optimization Algorithms; Particle Swarm Optimization Algorithms; Genetic Algorithms; Optimization Computing Algorithms; Neural Network Search & Large-Scale Optimization; Multi-objective Optimization.

reinforcement learning and stochastic optimization: Distributional Reinforcement Learning Marc G. Bellemare, Will Dabney, Mark Rowland, 2023-05-30 The first comprehensive guide to distributional reinforcement learning, providing a new mathematical formalism for thinking about decisions from a probabilistic perspective. Distributional reinforcement learning is a new mathematical formalism for thinking about decisions. Going beyond the common approach to reinforcement learning and expected values, it focuses on the total reward or return obtained as a consequence of an agent's choices—specifically, how this return behaves from a probabilistic perspective. In this first comprehensive guide to distributional reinforcement learning, Marc G. Bellemare, Will Dabney, and Mark Rowland, who spearheaded development of the field, present its key concepts and review some of its many applications. They demonstrate its power to account for many complex, interesting phenomena that arise from interactions with one's environment. The authors present core ideas from classical reinforcement learning to contextualize distributional topics and include mathematical proofs pertaining to major results discussed in the text. They guide the reader through a series of algorithmic and mathematical developments that, in turn, characterize, compute, estimate, and make decisions on the basis of the random return. Practitioners in disciplines as diverse as finance (risk management), computational neuroscience, computational psychiatry, psychology, macroeconomics, and robotics are already using distributional reinforcement learning, paving the way for its expanding applications in mathematical finance, engineering, and the life sciences. More than a mathematical approach, distributional reinforcement learning represents a new perspective on how intelligent agents make predictions and decisions.

reinforcement learning and stochastic optimization: Adaptive Agents and Multi-Agent Systems III. Adaptation and Multi-Agent Learning Karl Tuyls, Ann Nowe, Zahia Guessoum, Daniel Kudenko, 2008-02-08 This book contains selected and revised papers of the European Symposium on Adaptive and Learning Agents and Multi-Agent Systems (ALAMAS), editions 2005, 2006 and 2007, held in Paris, Brussels and Maastricht. The goal of the ALAMAS symposia, and this associated book, is to increase awareness and interest in adaptation and learning for single agents and mul- agent systems, and encourage collaboration between machine learning experts, softwareengineeringexperts, mathematicians, biologists and physicists, and give a representative overview of current state of affairs in this area. It is an inclusive forum where researchers can present recent work and discuss their newest ideas for a first time with their peers. The symposia series focuses on all aspects of adaptive and learning agents and multi-agent systems, with a particular emphasis on how to modify established learning techniques and/or create new learning paradigms to address the many challenges presented by complex real-world problems. These symposia were a great success and provided a forum for the presentation of new ideas and results bearing on the conception of adaptation and learning for single agents and multi-agent systems. Over these three editions we received 51 submissions, of which 17 were carefully selected, including one invited paper of this year's invited speaker Simon Parsons. This is a very competitive acceptance rate of approximately 31%, which, together with two review cycles, has led to a high-quality LNAI volume. We hope that our readers will be inspired by the papers included in this volume.

Related to reinforcement learning and stochastic optimization

Reinforcement - Wikipedia Reinforcement is an important component of operant conditioning and behavior modification. The concept has been applied in a variety of practical areas, including parenting, coaching,

REINFORCEMENT Definition & Meaning - Merriam-Webster The meaning of REINFORCEMENT is the action of strengthening or encouraging something : the state of being reinforced. How to use reinforcement in a sentence

Types of Reinforcement in Psychology: Definition and Examples Reinforcement strengthens behavior. Learn more about the reinforcement definition in psychology, along with examples and how it works to modify behavior

What Is Reinforcement? Psychology, Definition, And Applications Reinforcement psychology involves the use of providing something or taking it away to achieve a desired behavior. Primary reinforcement occurs naturally, while secondary

REINFORCEMENT | English meaning - Cambridge Dictionary REINFORCEMENT definition: 1. the act of making something stronger: 2. soldiers sent to join an army to make it stronger: 3. Learn more

Reinforce vs Reenforce: The Critical Mistake You Must Avoid! In summary, reinforce is the accepted spelling, and reinforcement is the correct noun form that represents the concept of providing strength or support. There's no need to use

Reinforcement Definition & Meaning | Britannica Dictionary REINFORCEMENT meaning: 1 : people and supplies that are sent to help or support an army, military force, etc.; 2 : the act of strengthening or encouraging something

Reinforcement - definition of reinforcement by The Free Dictionary 1. the act of reinforcing; the state of being reinforced. 2. something that reinforces or strengthens. 3. Often, reinforcements. an additional supply of personnel, ships, aircraft, etc., for a military

APA Dictionary of Psychology See also negative reinforcement; positive reinforcement; schedule of reinforcement. A trusted reference in the field of psychology, offering more than 25,000 clear

What Is Reinforcement in Operant Conditioning? - Verywell Mind Reinforcement is an important concept in operant conditioning and the learning process. Learn how it's used and see conditioned reinforcer examples in everyday life

Reinforcement - Wikipedia Reinforcement is an important component of operant conditioning and behavior modification. The concept has been applied in a variety of practical areas, including parenting, coaching,

REINFORCEMENT Definition & Meaning - Merriam-Webster The meaning of REINFORCEMENT is the action of strengthening or encouraging something : the state of being reinforced. How to use reinforcement in a sentence

Types of Reinforcement in Psychology: Definition and Examples Reinforcement strengthens behavior. Learn more about the reinforcement definition in psychology, along with examples and how it works to modify behavior

What Is Reinforcement? Psychology, Definition, And Applications Reinforcement psychology involves the use of providing something or taking it away to achieve a desired behavior. Primary reinforcement occurs naturally, while secondary

REINFORCEMENT | English meaning - Cambridge Dictionary REINFORCEMENT definition: 1. the act of making something stronger: 2. soldiers sent to join an army to make it stronger: 3. Learn more

Reinforce vs Reenforce: The Critical Mistake You Must Avoid! In summary, reinforce is the accepted spelling, and reinforcement is the correct noun form that represents the concept of providing strength or support. There's no need to use

Reinforcement Definition & Meaning | Britannica Dictionary REINFORCEMENT meaning: 1 : people and supplies that are sent to help or support an army, military force, etc.; 2 : the act of

strengthening or encouraging something

Reinforcement - definition of reinforcement by The Free Dictionary 1. the act of reinforcing; the state of being reinforced. 2. something that reinforces or strengthens. 3. Often, reinforcements. an additional supply of personnel, ships, aircraft, etc., for a military

APA Dictionary of Psychology See also negative reinforcement; positive reinforcement; schedule of reinforcement. A trusted reference in the field of psychology, offering more than 25,000 clear

What Is Reinforcement in Operant Conditioning? - Verywell Mind Reinforcement is an important concept in operant conditioning and the learning process. Learn how it's used and see conditioned reinforcer examples in everyday life

Reinforcement - Wikipedia Reinforcement is an important component of operant conditioning and behavior modification. The concept has been applied in a variety of practical areas, including parenting, coaching,

REINFORCEMENT Definition & Meaning - Merriam-Webster The meaning of REINFORCEMENT is the action of strengthening or encouraging something : the state of being reinforced. How to use reinforcement in a sentence

Types of Reinforcement in Psychology: Definition and Examples Reinforcement strengthens behavior. Learn more about the reinforcement definition in psychology, along with examples and how it works to modify behavior

What Is Reinforcement? Psychology, Definition, And Applications Reinforcement psychology involves the use of providing something or taking it away to achieve a desired behavior. Primary reinforcement occurs naturally, while secondary

REINFORCEMENT | English meaning - Cambridge Dictionary REINFORCEMENT definition: 1. the act of making something stronger: 2. soldiers sent to join an army to make it stronger: 3. Learn more

Reinforce vs Reenforce: The Critical Mistake You Must Avoid! In summary, reinforce is the accepted spelling, and reinforcement is the correct noun form that represents the concept of providing strength or support. There's no need to use

Reinforcement Definition & Meaning | Britannica Dictionary REINFORCEMENT meaning: 1 : people and supplies that are sent to help or support an army, military force, etc.; 2 : the act of strengthening or encouraging something

Reinforcement - definition of reinforcement by The Free Dictionary 1. the act of reinforcing; the state of being reinforced. 2. something that reinforces or strengthens. 3. Often, reinforcements. an additional supply of personnel, ships, aircraft, etc., for a military

APA Dictionary of Psychology See also negative reinforcement; positive reinforcement; schedule of reinforcement. A trusted reference in the field of psychology, offering more than 25,000 clear

What Is Reinforcement in Operant Conditioning? - Verywell Mind Reinforcement is an important concept in operant conditioning and the learning process. Learn how it's used and see conditioned reinforcer examples in everyday life

Reinforcement - Wikipedia Reinforcement is an important component of operant conditioning and behavior modification. The concept has been applied in a variety of practical areas, including parenting, coaching, therapy,

REINFORCEMENT Definition & Meaning - Merriam-Webster The meaning of REINFORCEMENT is the action of strengthening or encouraging something : the state of being reinforced. How to use reinforcement in a sentence

Types of Reinforcement in Psychology: Definition and Examples Reinforcement strengthens behavior. Learn more about the reinforcement definition in psychology, along with examples and how it works to modify behavior

What Is Reinforcement? Psychology, Definition, And Applications Reinforcement psychology involves the use of providing something or taking it away to achieve a desired behavior. Primary reinforcement occurs naturally, while secondary

REINFORCEMENT | English meaning - Cambridge Dictionary REINFORCEMENT definition: 1.

the act of making something stronger: 2. soldiers sent to join an army to make it stronger: 3. Learn more

Reinforce vs Reenforce: The Critical Mistake You Must Avoid! In summary, reinforce is the accepted spelling, and reinforcement is the correct noun form that represents the concept of providing strength or support. There's no need to use

Reinforcement Definition & Meaning | Britannica Dictionary REINFORCEMENT meaning: 1 : people and supplies that are sent to help or support an army, military force, etc.; 2 : the act of strengthening or encouraging something

Reinforcement - definition of reinforcement by The Free 1. the act of reinforcing; the state of being reinforced. 2. something that reinforces or strengthens. 3. Often, reinforcements. an additional supply of personnel, ships, aircraft, etc., for a military

APA Dictionary of Psychology See also negative reinforcement; positive reinforcement; schedule of reinforcement. A trusted reference in the field of psychology, offering more than 25,000 clear

What Is Reinforcement in Operant Conditioning? - Verywell Mind Reinforcement is an important concept in operant conditioning and the learning process. Learn how it's used and see conditioned reinforcer examples in everyday life

Reinforcement - Wikipedia Reinforcement is an important component of operant conditioning and behavior modification. The concept has been applied in a variety of practical areas, including parenting, coaching, therapy,

REINFORCEMENT Definition & Meaning - Merriam-Webster The meaning of REINFORCEMENT is the action of strengthening or encouraging something : the state of being reinforced. How to use reinforcement in a sentence

Types of Reinforcement in Psychology: Definition and Examples Reinforcement strengthens behavior. Learn more about the reinforcement definition in psychology, along with examples and how it works to modify behavior

What Is Reinforcement? Psychology, Definition, And Applications Reinforcement psychology involves the use of providing something or taking it away to achieve a desired behavior. Primary reinforcement occurs naturally, while secondary

REINFORCEMENT | English meaning - Cambridge Dictionary REINFORCEMENT definition: 1. the act of making something stronger: 2. soldiers sent to join an army to make it stronger: 3. Learn more

Reinforce vs Reenforce: The Critical Mistake You Must Avoid! In summary, reinforce is the accepted spelling, and reinforcement is the correct noun form that represents the concept of providing strength or support. There's no need to use

Reinforcement Definition & Meaning | Britannica Dictionary REINFORCEMENT meaning: 1 : people and supplies that are sent to help or support an army, military force, etc.; 2 : the act of strengthening or encouraging something

Reinforcement - definition of reinforcement by The Free 1. the act of reinforcing; the state of being reinforced. 2. something that reinforces or strengthens. 3. Often, reinforcements. an additional supply of personnel, ships, aircraft, etc., for a military

APA Dictionary of Psychology See also negative reinforcement; positive reinforcement; schedule of reinforcement. A trusted reference in the field of psychology, offering more than 25,000 clear

What Is Reinforcement in Operant Conditioning? - Verywell Mind Reinforcement is an important concept in operant conditioning and the learning process. Learn how it's used and see conditioned reinforcer examples in everyday life

Related to reinforcement learning and stochastic optimization

Stochastic Collusion and the Power Law of Learning: A General Reinforcement Learning Model of Cooperation (JSTOR Daily8y) Concerns about models of cultural adaptation as analogs of genetic selection have led cognitive game theorists to explore learning-theoretic specifications. Two prominent examples, the Bush-Mosteller

Stochastic Collusion and the Power Law of Learning: A General Reinforcement Learning Model of Cooperation (JSTOR Daily8y) Concerns about models of cultural adaptation as analogs of genetic selection have led cognitive game theorists to explore learning-theoretic specifications. Two prominent examples, the Bush-Mosteller

Reinforcement learning is making a buzz in space (MilitaryNews.com14d) A (NRL) research team successfully conducted the first reinforcement learning (RL) control of a free-flyer in space on May 27

Reinforcement learning is making a buzz in space (MilitaryNews.com14d) A (NRL) research team successfully conducted the first reinforcement learning (RL) control of a free-flyer in space on May 27

AI and Macroeconomic Modeling: Deep Reinforcement Learning in an RBC model (International Monetary Fund2y) Download PDF More Formats on IMF eLibrary Order a Print Copy Create Citation This study seeks to construct a basic reinforcement learning-based AI-macroeconomic simulator. We use a deep RL (DRL)

AI and Macroeconomic Modeling: Deep Reinforcement Learning in an RBC model (International Monetary Fund2y) Download PDF More Formats on IMF eLibrary Order a Print Copy Create Citation This study seeks to construct a basic reinforcement learning-based AI-macroeconomic simulator. We use a deep RL (DRL)

Related Articles

- [preguntas y respuestas del examen de manejo en oregon](#)
- [a new career at 50](#)
- [5b 3 el accidente de diego answers](#)

[Back to Home](#)