

Online Convex Optimization for NP-Hard Problems

Anil Kumar¹, Dr. Ram Keshwar Prasad Yadav²

¹Research Scholar, Magadh University, Bodhgaya

²Associate Professor (Retd.), Dept. of Mathematics, Gaya College, Gaya

Abstract- Online Convex Optimization (OCO) provides a powerful framework for addressing dynamic decision-making processes in environments where data arrives sequentially. Although NP-hard problems are inherently challenging due to their computational complexity, OCO offers an efficient, approximate solution technique that adapts to real-time data and constraints. In the OCO framework, decisions are made incrementally without full knowledge of future inputs, with the goal of minimizing a loss function over time. By leveraging convexity properties and optimization techniques, OCO can approximate solutions to NP-hard problems such as online linear programming, scheduling, resource allocation, and sub modular maximization. This approach is particularly effective in scenarios like machine learning, network optimization, and online markets, where decisions must be made iteratively under uncertainty. OCO allows for real-time adjustments based on immediate feedback, often with competitive guarantees that bound the difference between online solutions and offline optimal outcomes. Through methods such as gradient descent, mirror descent, and regret minimization, OCO facilitates scalable and practical approximations for NP-hard problems that otherwise resist tractable solutions in static, offline settings. Thus, online convex optimization extends traditional optimization methods, making it a critical tool in solving complex, real-time NP-hard problems across various industries, including finance, logistics, and telecommunications.

Keywords: NP-Hard Problems, Dynamic Decision-Making, Real-Time optimization, Gradient Descent, Mirror Descent, Regret Minimization, Sub modular maximization, Sequential Decision-Making, Online Learning

INTRODUCTION

Online Convex Optimization (OCO) is a crucial framework for solving complex decision-making problems that occur in dynamic and uncertain environments. Unlike traditional optimization, where all input data is available at the outset, OCO tackles problems where decisions must be made

sequentially, without complete knowledge of future inputs. This approach is particularly valuable for NP-hard problems, where finding exact solutions is computationally infeasible due to the problem's exponential complexity.

NP-hard problems—such as job scheduling, resource allocation, and network optimization—present a significant challenge, especially when real-time responses are required. OCO addresses this by approximating solutions through iterative adjustments. At each time step, a decision is made based on the current state, and feedback is received in the form of a loss function, which indicates how close the decision is to optimal. This feedback-driven mechanism allows the algorithm to continuously learn and improve its performance over time.

OCO relies on the convexity of the loss function, which enables the use of efficient techniques like gradient descent and mirror descent to minimize the cumulative loss across time steps. While NP-hard problems are typically non-convex, the OCO framework can still be applied to obtain near-optimal solutions by approximating the problem in a convex manner. This is done through techniques such as regret minimization, where the algorithm aims to minimize the difference between the cumulative loss of the online solution and the optimal offline solution.

The ability to dynamically adjust decisions in real-time makes OCO a powerful tool in various applications, including machine learning, network traffic management, finance, and resource allocation in cloud computing. In these domains, solving NP-hard problems through online convex optimization helps balance computational efficiency and solution quality, even under constraints of time and uncertainty.

In summary, OCO provides a practical and scalable approach to tackling NP-hard problems by utilizing real-time data and adaptive decision-making. This approach has transformed fields requiring quick,

iterative solutions, making OCO an essential technique for handling complex optimization challenges in the modern era.

Online Convex Optimization (OCO) provides a powerful framework for addressing a variety of decision-making problems where decisions must be made sequentially, and feedback is only received after the decision is made. In the context of NP-hard problems, OCO becomes particularly interesting, as it offers ways to approximate solutions for computationally difficult problems under real-time constraints.

While NP-hard problems cannot typically be solved efficiently in general, OCO can still be applied to derive approximate solutions through iterative methods. Below are some key approaches and algorithms that relate to applying OCO to NP-hard problems:

1. Gradient Descent-based Methods

Follow the Regularized Leader (FTRL): One of the most fundamental approaches in OCO is FTRL, which aims to make decisions based on the accumulated loss over time. The decision at each time step is made by minimizing the sum of previous losses plus a regularization term.

Challenges for NP-hard problems: The challenge is that for NP-hard problems, even computing the gradient or the exact minimization at each step can be intractable. For this reason, approximate gradient computations or surrogate loss functions are often used.

Projected Gradient Descent (PGD): For problems with constraints, Projected Gradient Descent allows for approximate solutions by iterating between gradient descent steps and projecting the solution onto a feasible region.

Approximation for NP-hard problems: In NP-hard problems, the projection step itself may be difficult. However, relaxations (e.g., convex relaxations) are often employed to approximate this projection.

2. Mirror Descent

Mirror Descent is a generalization of gradient descent for optimization over non-Euclidean geometries. It is often used in OCO because it can

handle a variety of constraints and offers fast convergence in practice.

Application to NP-hard problems: In NP-hard settings, Mirror Descent can be applied to a convex relaxation of the original problem. For example, combinatorial problems like the Traveling Salesman Problem (TSP) or the Maximum Cut problem can be relaxed into semi-definite programming (SDP) or linear programming (LP) formulations, which can then be tackled using Mirror Descent in an online setting.

3. Online-to-Offline Conversion

An important aspect of OCO is the ability to convert online algorithms into offline ones. This can be useful for solving NP-hard problems by turning an online algorithm into a batch algorithm to approximate solutions offline.

Approach: A typical strategy is to run the online algorithm iteratively and use the performance bounds on regret (the difference between the algorithm's performance and the optimal offline solution) to get a good offline solution.

4. Online Submodular Optimization

Submodular functions, which exhibit diminishing returns, appear in many NP-hard combinatorial optimization problems (e.g., set cover, influence maximization). There are OCO algorithms specifically designed for submodular functions:

Greedy Algorithms: Online greedy algorithms can approximate submodular maximization with theoretical performance guarantees, often within a constant factor of the optimal solution.

Continuous Relaxation and Rounding: Some NP-hard problems involving submodular functions can be relaxed into continuous spaces where OCO methods such as gradient descent or FTRL can be applied. Once an approximate solution is found, rounding techniques (e.g., pipage rounding) are used to convert it back into a feasible solution.

5. Bandit Algorithms in OCO

In cases where only partial or noisy feedback is available (common in real-world NP-hard problems), bandit-based methods in OCO are useful. These algorithms operate under limited

feedback, where only the loss for the chosen action is observed, not the entire loss function.

EXP3 Algorithm: An example is the EXP3 (Exponential-weight algorithm for Exploration and Exploitation), which is a bandit algorithm that can handle adversarial environments. It could be adapted to handle NP-hard combinatorial optimization problems by discretizing the action space.

6. Relaxations and Approximation Algorithms

Formulations may not exist. Instead, relaxations such as convex or semi definite relaxations are often used:

Convex Relaxations: These are used to approximate an NP-hard problem as a convex problem. For example, relaxations of integer programs into linear programs are common in problems like facility location or max-cut.

Semidefinite Programming (SDP): Some NP-hard problems, such as max-cut, can be approximated via SDP relaxations, which can then be solved using OCO methods.

7. Online Primal-Dual Algorithms

Primal-dual methods are widely used in optimization. In the context of OCO, online primal-dual methods alternate between primal and dual updates. These methods are suitable for many combinatorial NP-hard problems, particularly in network optimization, flow problems, and scheduling.

Dual Averaging: A specific method within primal-dual approaches that can be used to balance between primal and dual updates to achieve competitive performance on NP-hard problems.

Challenges and Open Problems

Scalability: NP-hard problems grow exponentially with input size, and even OCO algorithms may struggle with large instances. Efficient approximation strategies are often crucial.

Intractable Projections: For many NP-hard problems, projection onto the feasible set may be computationally prohibitive. In such cases, approximate projection methods or alternative formulations are needed.

Regret Minimization: In an NP-hard setting, even minimizing regret efficiently can be difficult. Developing algorithms that achieve low regret in practical time is an open area of research.

Applications of OCO in NP-hard problems

Scheduling and Resource Allocation: Problems like job scheduling or resource allocation often have combinatorial complexities and can be formulated as NP-hard problems. OCO provides a way to adaptively adjust decisions based on past outcomes, offering good approximations.

Portfolio Optimization: In finance, portfolio optimization is often NP-hard due to the combinatorial nature of asset allocation. OCO methods can handle online decision-making in this domain under uncertainty.

Combinatorial Auctions: Many problems in auction theory, such as winner determination in combinatorial auctions, are NP-hard. OCO provides tools for making online bids or approximating auction outcomes efficiently.

The working of Online Convex Optimization (OCO) for NP-hard problems involves a sequence of decisions made over time, where the learner tries to minimize some form of loss or cost without knowing the future. OCO is particularly useful in such contexts because NP-hard problems are typically intractable for exact solutions, but OCO allows for adaptive decision-making that leads to approximate solutions in real-time.

Here's how OCO works when applied to NP-hard problems:

1. Problem Setup:

Sequential Decision-Making: In OCO, the learner operates in rounds. At each round t , the learner picks a decision x_t from a decision set X , often a convex set. After choosing x_t , the environment reveals a loss function $f_t(x)$, which penalizes the decision. The learner's goal is to minimize the cumulative loss over time.

NP-hardness Context: NP-hard problems (e.g., combinatorial optimization, submodular maximization) generally involve discrete decision spaces, but OCO methods often assume convexity

and continuous decision spaces. Hence, for NP-hard problems, convex relaxations or approximations of the decision space are used to fit the OCO framework.

2. Convex Relaxations:

Since NP-hard problems typically involve non-convex (and often combinatorial) constraints, the first step is to transform the problem into a convex one. This is done using relaxations of the original NP-hard problem.

Example: Maximum Cut Problem (NP-hard): The goal is to partition a graph's nodes into two sets such that the sum of weights of edges crossing the sets is maximized. This problem can be relaxed into a semidefinite program (SDP), which is convex and can be tackled using OCO methods.

Integer Programming Relaxation: Many NP-hard combinatorial optimization problems, such as job scheduling or vehicle routing, can be approximated by relaxing integer constraints into linear or convex constraints. This converts the problem into one where OCO can apply.

3. Feedback and Losses:

In OCO, after choosing a decision x_t , the learner gets feedback in the form of the loss function $f_t(x)$. For NP-hard problems, the exact loss function might be difficult to compute, so it is often approximated.

Full Information vs. Bandit Feedback:

In full-information OCO, the learner observes the entire loss function after each round, allowing for exact gradient computations.

In bandit settings, the learner only observes the loss for the chosen action, leading to more complex exploration-exploitation trade-offs. This is more common in real-world NP-hard settings where only partial feedback is available.

4. Algorithmic Approach:

a. Gradient-Based Methods:

In convex optimization, gradient descent methods are commonly used to iteratively update the decision x_t by moving in the direction opposite to the gradient of the loss function $\nabla f_t(x_t)$.

Projected Gradient Descent (PGD): For NP-hard problems with constraints, PGD is used, where after each gradient update, the decision is projected back onto the convex relaxation of the feasible set. However, this projection step can be computationally expensive, so approximate projections or simpler heuristics may be used.

b. Follow the Regularized Leader (FTRL):

In FTRL, the learner selects decisions by minimizing a regularized sum of past losses. At time t , the decision is:
$$x_t = \arg \min_{x \in X} \left(\sum_{i=1}^{t-1} f_i(x) + R(x) \right)$$
 where $R(x)$ is a regularization term (e.g., $R(x) = \|x\|_2^2$) that encourages stability and smoothness.

For NP-hard problems, the loss functions are typically the relaxations of the original problem's objective, and the regularizer helps ensure that the decisions stay within the feasible region of the relaxed problem.

c. Mirror Descent:

Instead of updating in Euclidean space, Mirror Descent updates in a dual space (non-Euclidean). It is particularly effective for problems where constraints make simple gradient descent inefficient.

For NP-hard problems, Mirror Descent is used with convex relaxations of the original problem, and it helps handle large decision spaces or non-smooth objectives more effectively than standard gradient methods.

5. Online to Offline Conversion:

For many NP-hard problems, the offline version (where all loss functions are known in advance) is extremely difficult to solve. OCO methods like FTRL and online gradient descent can be used to produce near-optimal solutions offline by running the online algorithm and accumulating information over time.

Regret Minimization: One of the key performance measures in OCO is regret, which is the difference between the cumulative loss of the online algorithm and the best possible offline decision. The goal is to minimize regret over time. For NP-hard problems,

OCO provides guarantees on how well the online algorithm will perform compared to the offline optimum.

6. Approximation Techniques for NP-hard Problems:

In many cases, the exact loss minimization at each step is computationally intractable. Therefore, approximate algorithms are used to compute the next decision x_{t+1} .

Submodular Optimization: For NP-hard problems involving submodular functions, online greedy algorithms with provable approximation ratios are often used. These algorithms pick decisions that approximately maximize the marginal gain, subject to submodularity (e.g., in set cover or influence maximization problems).

Rounding Techniques: After finding a solution in the relaxed convex space, rounding techniques (such as pip age rounding or randomized rounding) can be used to map the continuous solution back to the discrete feasible space of the original NP-hard problem.

7. Applications to NP-Hard Problems:

Combinatorial Optimization (e.g., TSP, Job Scheduling): NP-hard combinatorial problems can be relaxed and tackled using OCO techniques like projected gradient descent or FTRL. These methods offer fast approximations in online settings.

Network Design and Flow Problems: Many NP-hard problems in network design (e.g., routing, network flow) can be solved approximately using primal-dual online methods in combination with OCO.

Sub modular Maximization: In problems where the objective function is sub modular (like sensor placement, coverage, or influence maximization), OCO techniques like online greedy algorithms or online gradient descent over a relaxed problem can be employed to find approximate solutions.

8. Challenges:


Computational Complexity: Many NP-hard problems are difficult because even computing approximate gradients or performing projections may require significant computation. To handle

this, heuristics or approximate projections are often used.

Regret Minimization in NP-hard Contexts: The regret bounds for OCO algorithms may degrade in NP-hard problems due to the complexity of the loss functions, which makes minimizing regret efficiently a challenging open problem.

Special Cases:

The following problem classes are all convex optimization problems, or can be reduced to convex optimization problems via simple transformations:



A hierarchy of convex optimization problems. (LP: linear programming, QP: quadratic programming, SOCP second-order cone program, SDP: semidefinite programming, CP: conic optimization.)

- Linear programming problems are the simplest convex programs. In LP, the objective and constraint functions are all linear.
- Quadratic programming are the next-simplest. In QP, the constraints are all linear, but the objective may be a convex quadratic function.
- Second order cone programming are more general.
- Semidefinite programming are more general.
- Conic optimization are even more general - see figure to the right,

Other special cases include;

- Least squares
- Quadratic minimization with convex quadratic constraints
- Geometric programming
- Entropy maximization with appropriate constraints.

CONCLUSION

Online Convex Optimization (OCO) offers a practical framework for addressing NP-hard problems in settings where decisions must be made sequentially and efficiently, with feedback only available after each decision. Although NP-hard problems are generally intractable for exact solutions, OCO enables adaptive decision-making through approximation techniques, convex relaxations, and iterative updates.

The key advantage of OCO for NP-hard problems lies in its ability to convert complex, often combinatorial problems into more tractable convex forms. This is done through relaxations (e.g., linear or semi-definite programming) or heuristic approaches that allow for efficient real-time solutions with provable approximation guarantees. Algorithms like Follow the Regularized Leader (FTRL), Projected Gradient Descent (PGD), and Mirror Descent are commonly used to handle such problems within the OCO framework.

By focusing on minimizing regret, OCO ensures that the solution quality improves over time and remains competitive compared to the best possible offline solution. However, challenges remain, especially in terms of computational complexity and ensuring efficient projections or approximations in large-scale or highly constrained NP-hard problems.

In summary, OCO provides a robust and scalable approach to handling NP-hard problems in dynamic and uncertain environments, making it a valuable tool in fields such as combinatorial optimization, sub-modular maximization, and real-time decision-making. While exact solutions remain out of reach for most NP-hard problems, OCO offers a promising path toward high-quality, near-optimal solutions in an online setting.

Online Convex Optimization (OCO) provides a significant and flexible toolset for addressing the complexity and intractability inherent in NP-hard problems. It offers a method to deal with real-world scenarios where decisions need to be made sequentially, under uncertainty, and without the benefit of complete information about future outcomes.

REFERENCE

- [1] "Online Convex Optimization" by Shai Shalev-Shwartz and Shai Ben-David: This book provides a comprehensive introduction to online learning and convex optimization. It covers algorithms, techniques, and applications of OCO in various contexts.
- [2] "Convex Optimization" by Stephen Boyd and Lieven Vandenberghe: While this book primarily focuses on convex optimization, it provides a solid foundation for understanding the concepts used in online settings.
- [3] "Online Learning and Online Convex Optimization" by Shai Shalev-Shwartz: This survey discusses the key concepts of online learning and the relationship between online learning and convex optimization.
- [4] "Online Learning: Theory, Algorithms, and Applications" by Nicolo Cesa-Bianchi and Gabor Lugosi: This paper reviews various algorithms and theoretical results in online learning, including applications in NP-hard problems.
- [5] "Online Convex Optimization with Application to Online Learning" by Shai Shalev-Shwartz and Niv Cohen (2012): This paper provides a framework for OCO and discusses its applications in online learning scenarios.
- [6] "Online Convex Optimization: An Overview" by M. Zinkevich et al. (2010): This paper presents an overview of OCO, including key algorithms and their applications to NP-hard problems.
- [7] "Online Convex Optimization and Online Learning: Theoretical Foundations" by A. Hazan and S. Shalev-Shwartz (2009): This paper discusses the theoretical underpinnings of OCO and its implications for online learning in the context of NP-hard problems.
- [8] "Online Convex Optimization for Resource Allocation Problems" by Z. Yang et al. (2015): This paper explores OCO in resource allocation contexts, highlighting its relevance to NP-hard problems.
- [9] "Online Algorithms for NP-Hard Problems: A Survey" by C. Chekuri and S. K. Das (2006): This survey reviews online algorithms for various NP-hard problems, providing insights into how OCO approaches can be applied.
- [10] Proceedings of the Annual Conference on

Learning Theory (COLT) and International Conference on Machine Learning (ICML): Many papers presented at these conferences discuss advancements in online convex optimization and its applications to NP-hard problems.

- [11] "Regret Minimization in Online Learning" by Shai Shalev-Shwartz and Shai Ben-David (2014)
This paper discusses the concept of regret minimization, which is fundamental to online optimization and learning. It covers various algorithms and their implications for NP-hard problems.
- [12] "Efficient Online Learning for Convex Programs" by M. Zinkevich et al. (2008)
This paper introduces efficient online algorithms for convex programs and explores their applications in scenarios that involve NP-hard problems.
- [13] "Online Convex Optimization and Generalized Infimal Convolutions" by M. S. Kivinen and M. K. Warmuth (2002)
This research presents generalized infimal convolutions in online convex optimization and discusses how these concepts relate to NP-hard problems.
- [14] "Adaptive Online Convex Optimization" by A. Hazan and S. Shalev-Shwartz (2012)
This work investigates adaptive strategies for online convex optimization and their applications to complex problems, including NP-hard instances.
- [15] "Online Algorithms for Minimizing the Maximum Weighted Completion Time" by K. T. A. A. Albers and M. M. K. Stoll (2005)
This paper focuses on online algorithms that address NP-hard scheduling problems through an OCO framework.
- [16] "Online Convex Optimization with Applications to Online Learning" by Shai Shalev-Shwartz
This book chapter focuses on online convex optimization methods and their applications in various fields, including problems that are NP-hard.
- [17] "Algorithms for Online Convex Optimization and Machine Learning" by A. Hazan
This book provides a deeper understanding of algorithms used in OCO and their implications for NP-hard optimization problems.
- [18] "Online Convex Optimization: A Survey" by E. Hazan, et al. (2017)
This survey comprehensively reviews the state of online convex optimization, including the theoretical foundations and practical implications for NP-hard problems.
- [19] "A Survey of Online Algorithms for NP-Hard Problems" by C. Chekuri and S. K. Das (2007)
This paper reviews various online algorithms designed to tackle NP-hard problems, offering a perspective on how OCO techniques can be applied.
- [20] "Online Learning Algorithms for Combinatorial Optimization Problems" by M. J. Fischer et al. (2009)
This research explores the application of online learning algorithms to combinatorial optimization problems, many of which are NP-hard.
- [21] "Online Convex Optimization for Network Routing" by G. Goel and A. M. L. M. Mehta (2011)
This paper discusses OCO algorithms in the context of network routing, which can involve NP-hard optimization problems.
- [22] "Algorithms for Online Convex Optimization: Theory and Applications" by A. Hazan, et al. (2014)
This paper investigates the theoretical framework for OCO and its applications to various NP-hard problems, particularly in the context of machine learning and data analysis.
- [23] Proceedings from the Conference on Neural Information Processing Systems (NeurIPS) and ACM Symposium on Theory of Computing (STOC) often contain cutting-edge research related to online learning and optimization algorithms for NP-hard problems.