Online Convex Optimization

Xinlei Yi, Laboratory for Information & Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, United States
Tao Yang, State Key Laboratory of Synthetical Automation for Process Industries, Northeastern University, Shenyang, China
Karl H. Johansson, Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden

Abstract

Online convex optimization is a sequential decision-making problem with a sequence of arbitrarily varying convex loss functions. This problem has gained renewed interest since it is a promising framework for machine learning and has wide applications. This chapter provides an overview of online convex optimization and its variations with long-term and time-varying constraints. To address critical issues in centralized processing, such as data privacy, data security, and single point failures, distributed online convex optimization and its variations are also reviewed. The differences among these variations are briefly discussed and state-of-the-art results in each category are summarized.

Key Points

- Introduce online convex optimization and its variations with long-term and time-varying constraints.
- Discuss constraint violation metrics.
- Introduce distributed online convex optimization and its variations.
- Present challenges when analyzing distributed online convex optimization algorithms.

Introduction

Online convex optimization is a promising framework for machine learning and has wide applications such as online classification (Crammer *et al.*, 2006), dictionary learning (Mairal *et al.*, 2009), and online advertising (Goldfarb and Tucker, 2011). It can be traced back at least to the 1990's (Cesa-Bianchi *et al.*, 1996). Simply speaking, online convex optimization is a sequential decision-making problem with a sequence of arbitrarily varying convex loss functions. At each round, a decision maker selects a decision from the decision set and then the loss function at this round is revealed. The goal of the decision maker is to minimize the loss accumulated over time. For an online convex optimization algorithm, the standard performance metric is regret, which is the performance gap between the decision sequence induced by the algorithm and a benchmark in hindsight. In other words, as pointed out in Hazan (2016), the goal of online algorithms is to minimize regret, rather than the optimization error (which is ill-defined in an online setting). Note that at each round the loss function can be arbitrarily chosen by the adversary, especially with no probabilistic model imposed on the choices. This is the key difference between online and stochastic convex optimization.

Distributed optimization methods are core to important applications because of their flexibility and scalability to large-scale datasets and systems, and ability to handle data privacy and locality constraints (Koloskova *et al.*, 2019). Motivated by this, distributed online convex optimization has also been extensively studied. In this problem, there are a group of agents (decision makers). At each round, each agent selects a decision, and then a portion of the global loss function is revealed to this agent only. The goal of the agents is to minimize a network-wide accumulated loss, and the corresponding performance measure is network regret, i.e., the average of all individual regrets. Each agent's individual regret is the difference between the cumulative global losses evaluated at this agent's decision sequence and a benchmark in hindsight. As a result, under network regret, any local decision sequence can be adopted by the network as a global decision sequence. However, due to partial knowledge of the global loss function for each agent, agents need to collaborate, which is the main difference between centralized and distributed online convex optimization. Therefore, it is needed to bound the error between agents' decisions when analyzing network regret bound, which is the key challenge when studying distributed online convex optimization.

The outline of this chapter is as follows. Section "Centralized Settings" presents the online convex optimization problem and discusses how to deal with long-time and time-varying constraints. Distributed online convex optimization is introduced in Section "Distributed Settings" together with constraints and some variations. Conclusions and some open problems are given in Section "Conclusions".

Centralized Settings

Online convex optimization can be understood as a repeated game between a learner and an adversary (Shalev-Shwartz, 2012). At round t of the game, the learner selects a point x_t from a known closed convex set $\mathcal{X} \subseteq \mathbb{R}^p$, and the adversary arbitrarily chooses a convex loss function $f_t: \mathbb{R}^p \to \mathbb{R}$. After that, the loss function f_t is revealed to the learner who suffers a loss $f_t(x_t)$. The goal of the learner is to choose a sequence $x_{[T]} = (x_1, \ldots, x_T)$ such that her regret

$$\mathrm{Reg}(x_{[T]}, y_{[T]}) := \sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t(y_t)$$

is minimized, where T is the total number of rounds and $y_{[T]} = (y_1, \ldots, y_T)$ is a benchmark. In the literature, there are two commonly used benchmarks. One is the optimal dynamic decision sequence $y_{[T]} = x^*_{[T]} = (x^*_1, \ldots, x^*_T)$ solving the following constrained convex optimization problem when the sequence of loss functions is known a priori:

$$\min_{x_1, \ldots, x_T \in \mathcal{X}} \sum_{t=1}^{T} f_t(x_t).$$

In this case, $\mathrm{Reg}(x_{[T]}, x^*_{[T]})$ is called the dynamic regret. The other benchmark is the optimal static decision sequence $y_{[T]} = \check{x}^*_{[T]} = (\check{x}^*_T, \ldots, \check{x}^*_T)$, where $\check{x}^*_T$ is the optimal static decision solving

$$\min_{x \in \mathcal{X}} \sum_{t=1}^{T} f_t(x).$$

In this case, $\mathrm{Reg}(x_{[T]}, \check{x}^*_{[T]})$ is called the static regret. In online convex optimization, we are usually interested in finding an upper bound on the worst regret of an algorithm. Intuitively, an algorithm performs well if its static regret is sublinear as a function of T, i.e., the static regret bound has an order of $\mathcal{O}(T^\kappa)$ with κ being a constant in the interval $(0,1)$, since this implies that on average the algorithm performs as well as the best fixed strategy in hindsight as T goes to infinity.

It is known that the simple and popular projection-based online gradient descent algorithm

$$x_{t+1} = \mathcal{P}_{\mathcal{X}}(x_t - \alpha_t \partial f_t(x_t)), \tag{1}$$

where $\mathcal{P}_{\mathcal{X}}(\cdot)$ is the projection onto the closed convex set $\mathcal{X}$, $\alpha_t > 0$ is the stepsize and $\partial f_t(x_t)$ denotes the subgradient of f_t at x_t, achieves an $\mathcal{O}(\sqrt{T})$ static regret bound for loss functions with bounded subgradients (Zinkevich, 2003), i.e.,

$$\mathrm{Reg}(x_{[T]}, \check{x}^*_{[T]}) = \mathcal{O}(\sqrt{T}).$$

It was later shown that $\mathcal{O}(\sqrt{T})$ is a tight bound up to constant factors, and the static regret bound can be reduced under more stringent strong convexity conditions on the loss functions (Hazan *et al.*, 2007). When the set $\mathcal{X}$ is bounded, it was also shown in Zinkevich (2003) that the algorithm (1) achieves the following regret bound

$$\mathrm{Reg}(x_{[T]}, y_{[T]}) = \mathcal{O}(\sqrt{T}(1 + P_T)),$$

where

$$P_T = \sum_{t=1}^{T-1} \|y_{t+1} - y_t\|$$

is the path-length of the benchmark $y_{[T]}$.

By running the projection-based online gradient descent algorithm (1) $\mathcal{O}(log(T))$ times in parallel with different stepsizes and choosing the smallest regret through an expert-tracking algorithm, the optimal dynamic regret bound

$$\mathrm{Reg}(x_{[T]}, y_{[T]}) = \mathcal{O}(\sqrt{T(1 + P_T)})$$

was achieved in Zhang *et al.* (2018). The curvature of loss functions, such as strong convexity and smoothness, can be used to further reduce the dynamic regret bound (Zhao and Zhang, 2021).

Long-Term Constraints

Despite the simplicity of algorithm (1), its computational cost is crucial. The projection $\mathcal{P}_{\mathcal{X}}(\cdot)$ is easy to compute and even has a closed form solution when $\mathcal{X}$ is a simple set, e.g., a box or a ball. However, in practice, the constraint set $\mathcal{X}$ is often complex. For example, if $\mathcal{X}$ is characterized by inequalities as

$$\mathcal{X} = \{x: g(x) \le \mathbf{0}_m, \quad x \in \mathbb{X}\}, \tag{2}$$

where $\mathbb{X} \subseteq \mathbb{R}^p$ is a closed convex set and $g(x) = (g_1(x), ..., g_m(x))^T$ with each $g_i: \mathbb{R}^p \to \mathbb{R}$ being a convex function, then the projection $\mathcal{P}_{\mathcal{X}}(\cdot)$ yields a heavy computational burden. To tackle this challenge, online convex optimization with long-term constraints was considered in Mahdavi *et al.* (2012). In this case, instead of requiring $g(x_t) \le \mathbf{0}_m$ at each round, the constraint should only be satisfied in the long run. More specifically, the constraint violation

$$\left\| \left[\sum_{t=1}^{T} g(x_t) \right]_+ \right\| \tag{3}$$

should grow sublinearly, where $[\cdot]_+$ is the projection onto the nonnegative space. In other words, the learner is allowed sometimes to make decisions that do not belong to the set $\mathcal{X}$, but the overall sequence of chosen decisions must obey the constraint at the end by a vanishing convergence rate. This problem is normally solved by online primal–dual algorithms. For example, Yu and Neely (2020) achieved an $\mathcal{O}(\sqrt{T})$ static regret bound and an $\mathcal{O}(T^{1/4})$ constraint violation bound, and the bound for constraint violation was reduced to $\mathcal{O}(1)$ under Slater's condition, i.e., there exists a point $x_s \in \mathbb{X}$ and a constant $\epsilon_s > 0$, such that $g(x_s) \le -\epsilon_s \mathbf{1}_m$.

The constraint violation metric defined in (3) allows inequality constraint violations at many rounds as long as they are compensated by a strictly feasible constraint that has a large margin, since it takes the summation over rounds before the projection operation $[\cdot]_+$. In this way, although the constraint violation grows sublinearly, the constraints may not be satisfied at many time instances, which will restrict the theoretical results only to applications where the constraints have cumulative nature. This motivates researchers to consider stricter forms of constraint violation metric. Yuan and Lamperski (2018) considered cumulative constraint violation

$$\left\| \sum_{t=1}^{T} [g(x_t)]_+ \right\|, \tag{4}$$

and cumulative squared constraint violation

$$\sum_{t=1}^{T} \| [g(x_t)]_+ \|^2. \tag{5}$$

Both forms of metrics (4) and (5) take into account all constraints that are not satisfied, and the metric (4) is stricter than the constraint violation metric (3). In Yuan and Lamperski (2018), an $\mathcal{O}(T^{\max\{c,1-c\}})$ static regret bound and an $\mathcal{O}(T^{1-c/2})$ cumulative constraint violation bound were achieved, where $c \in (0,1)$ is a user-defined trade-off parameter enabling the trade-off between these two bounds. The cumulative constraint violation bound was reduced to $\mathcal{O}(T^{(1-c)/2})$ in Yi *et al.* (2021), and by using expert-tracking technique the optimal $\mathcal{O}(\sqrt{T(1 + P_T)})$ dynamic regret and an $\mathcal{O}(\sqrt{T})$ cumulative constraint violation were also achieved. Even without Slater's condition, the cumulative constraint violation bound was further reduced in Guo *et al.* (2022) by rectifying penalty imposed when making decisions. Moreover, when the loss functions are strongly convex, static regret bound can be reduced. The key idea to consider the stricter metrics (4) and (5) is to use the clipped constraint function $[g]_+$ to replace the original constraint function g.

Time-Varying Constraints

The problem above has been further extended to the time-varying constraints setting, i.e., the constraint function can be arbitrarily and adversarially designed, and is revealed to the decision maker after selecting its decision at each round. In this case, let g_t: $\mathbb{R}^p \to \mathbb{R}^m$ denote the convex constraint function of the t-th round and then the corresponding metric is constraint violation $\| [\sum_{t=1}^T g_t(x_t)]_+ \|$ or cumulative constraint violation $\| \sum_{t=1}^T [g_t(x_t)]_+ \|$. Online convex optimization with time-varying constraints has also been extensively studied in the literature, which usually proposed online primal-dual algorithms and achieved sublinear static regret and (cumulative) constraint violation bounds. For example, Sun *et al.* (2017) achieved an $\mathcal{O}(\sqrt{T})$ static regret bound and an $\mathcal{O}(T^{3/4})$ constraint violation bound. In Yu *et al.* (2017) and its online version, the bound for constraint violation was reduced to $\mathcal{O}(\sqrt{T})$ under Slater's condition. Guo *et al.* (2022) achieved an $\mathcal{O}(\sqrt{T})$ static regret bound and an $\mathcal{O}(T^{3/4})$ cumulative constraint violation bound. Compared to the bounds achieved in the long-term constraints setting, the (cumulative) constraint violation bounds achieved in the time-varying constraints setting are much larger since for the former case the constraint function is fixed and known a priori, while for the later case the constraint function is time-varying and known a posteriori. Moreover, knowing the fixed constraint function in advance is also the key reason why an $\mathcal{O}(1)$ cumulative constraint violation bound is achievable even without Slater's condition as shown in Guo *et al.* (2022).

Note that Slater's condition is an example of a constraint qualification that guarantees strong duality in convex optimization problems with inequality constraints (Boyd and Vandenberghe, 2004). This condition guarantees that the optimal solution of the primal problem can be achieved without violating any of the constraints. In the context of online convex optimization, under Slater's condition, algorithms can adaptively adjust the decision at each round to ensure that the constraints are not violated significantly, even if the optimization problem is evolving dynamically. Therefore, under Slater's condition it is expected that the (cumulative) constraint violation bounds for the time-varying constraints should be reduced. As mentioned before, smaller constraint violation bound was achieved in Yu *et al.* (2017) under Slater's condition. It is challenging to show that smaller cumulative constraint violation bound can be achieved under Slater's condition. The reason is as follows. To consider cumulative constraint violation, the clipped constraint function is used to replace the original constraint function. However, this becomes ineffective when also considering Slater's condition, as the use of the clipping operation renders Slater's condition ineffective. More specifically, the clipped function is nonnegative, and thus it is impossible to find a point where the value of the clipped function is strictly less than zero.

Distributed Settings

Distributed online convex optimization has been extensively studied, e.g., Mateos-Núnez and Cortés (2014) and Li *et al.* (2023). In distributed online convex optimization, there are a group of agents (decision makers). At each round t, each agent i selects a decision $x_{i,t} \in \mathcal{X}$, and after the selection the convex local loss function $f_{i,t}$: $\mathbb{R}^p \to \mathbb{R}$, a portion of the global loss function $f_t(x) = \frac{1}{n}\sum_{i=1}^n f_{i,t}(x)$, is revealed to agent i only. The goal of the agents is to minimize the network-wide accumulated loss, and the corresponding performance measure is network regret

$$\text{Net-Reg}(\{x_{i,t}\}, y_{[T]}) := \frac{1}{n}\sum_{i=1}^n \sum_{t=1}^T f_t(x_{i,t}) - \sum_{t=1}^T f_t(y_t).$$

Clearly, network regret is the average of all individual regrets $\{\text{Reg}(x_{i,[T]}, y_{[T]})\}$, where $x_{i,[T]} = (x_{i,1}, \ldots, x_{i,T})$ is each agent's decision sequence. The main reason for evaluating f_t at $x_{i,t}$ is to show that any local decision sequence $x_{i,[T]}$ has guaranteed performance and can thus be adopted by the network as a global decision sequence. If we only evaluate $f_{i,t}$ at $x_{i,t}$, the distributed problem becomes n independent online convex optimization problems since each agent does not need to consider other agents' loss functions.

Note that each agent alone cannot compute network regret since it does not know other agents' local loss functions. Agents can use a consensus protocol to collaborate. Therefore, agents need to communicate with each other. It is assumed that agents are allowed to share their decisions through a communication network modeled by a graph. This graph could be undirected or directed, fixed or time-varying, deterministic or stochastic, connected or jointly connected. Under mild assumptions on the graph, distributed online algorithms can achieve the same bounds for network regret as their centralized counterparts. For example, the distributed projection-based online gradient descent algorithm

$$x_{i,t+1} = \mathcal{P}_{\mathcal{X}}\left(\sum_{j=1}^n [W_t]_{ij} x_{j,t} - \alpha_t \partial f_{i,t}(x_{i,t})\right),$$

where $[W_t]_{ij}$ is the i-th row and j-th column element of the mixing matrix corresponding to the time-varying graph, achieves the optimal $\mathcal{O}(\sqrt{T})$ static network regret bound when local loss functions has bounded subgradients. Compared to the centralized scenario, the analysis of distributed online algorithms is more complicated since the consensus error $\|x_{i,t} - x_{j,t}\|$ needs to be analyzed due to the absence of $f_{j,t}(x_{i,t})$. The distributed online algorithm would be more complicated and the corresponding analysis would be more challenging if the graph is directed, e.g., Akbari *et al.* (2017).

Long-Term Constraints

Similar to the centralized scenario, in order to avoid the potential computation and/or storage challenge caused by the projection operator when using projection-based algorithms, distributed online convex optimization with long-term constraints has also been considered, e.g., Yuan *et al.* (2022). In this problem, the constraint set $\mathcal{X}$ is characterized by inequality constraints as described in (2), and each agent knows the simple set $\mathbb{X}$ and the constraint function g in advance. Similar to the centralized case, the decisions are selected from $\mathbb{X}$ instead of $\mathcal{X}$ and the inequality constraints should be satisfied in the long-term on average, which is measured by network constraint violation $\frac{1}{n}\sum_{i=1}^{n}\left\|[\sum_{t=1}^{T}g(x_{i,t})]_{+}\right\|$ or network cumulative constraint violation $\frac{1}{n}\sum_{i=1}^{n}\sum_{t=1}^{T}\|[g(x_{i,t})]_{+}\|$. The static network regret and constraint violation bounds achieved for distributed online convex optimization with long-term constraints are similar to their centralized counterparts. For example, Yuan *et al.* (2022) achieved an $\mathcal{O}(T^{\max\{c,1-c\}})$ static network regret bound and an $\mathcal{O}(T^{1-c/2})$ network cumulative constraint violation bound, which are the same as the bounds achieved in Yuan and Lamperski (2018), but the bound for network cumulative constraint violation is greater than that in Yi *et al.* (2021) and Guo *et al.* (2022). It is still an unsolved problem to achieve the same network (cumulative) constraint violation bound for distributed online convex optimization with long-term constraints as its centralized counterpart.

Time-Varying Constraints

The distributed online convex optimization with long-term constraints setting was extended to a more general scenario in Yi *et al.* (2023a), where the constraint function is time-varying and at each round the convex local constraint function $g_{i,t}: \mathbb{R}^{p} \to \mathbb{R}^{m_i}$, a coordinate block of the global constraint function $g_t(x) = [[g_{1,t}(x)]^T, \ldots, [g_{n,t}(x)]^T]^T$, is privately revealed to each agent after selecting its decision. In this case, the corresponding metric is network constraint violation $\frac{1}{n}\sum_{i=1}^{n}\left\|[\sum_{t=1}^{T}g_t(x_{i,t})]_{+}\right\|$ or network cumulative constraint violation $\frac{1}{n}\sum_{i=1}^{n}\sum_{t=1}^{T}\left\|[g_t(x_{i,t})]_{+}\right\|$. The key difference between distributed online convex optimization with long-term and time-varying constraints is that for the former case the constraint function is fixed and known a priori to each agent, while for the later case the constraint function is time-varying and only a coordinate block known a posteriori to each agent. Notably, the distributed online convex optimization with time-varying constraints encompasses its centralized counterpart and distributed online convex optimization with long-term constraints as special cases. Despite this, Yi *et al.* (2023a) managed to attain the same bounds for network regret and cumulative constraint violation as those achieved in Yuan *et al.* (2022) and Guo *et al.* (2022).

The challenging scenario where Slater's condition holds was then considered in Yi *et al.* (2023b). The main challenge is to consider cumulative constraint violation and Slater's condition at the same time. As mentioned before, to consider cumulative constraint violation, previous papers all replace the original constraint function with the clipped constraint function, however, the clipping operation renders Slater's condition ineffective. To tackle this problem, Yi *et al.* (2023b) proposed a distributed online primal-dual algorithm, which updates the dual variables by directly maximizing the regularized Lagrangian function and thus can be explicitly calculated using the clipped constraint function. Consequently, for the scenario without Slater's condition, all squared constraint violations can be accumulated, leveraging the fact that $a^T[a]_{+} = \|[a]_{+}\|^2$ for any vector a, and state-of-the-art (network) cumulative constraint violation bounds can be achieved. Moreover, for the scenario with Slater's condition, all constraint violations can be accumulated directly by summing the dual variables. As a result, Yi *et al.* (2023b) demonstrated that smaller (network) cumulative constraint violation bounds can be achieved under Slater's condition.

Other Variations

Other forms of distributed variation of the centralized online convex optimization have also been considered. For example, distributed online convex optimization with nonseparable global objectives was considered in Lee *et al.* (2018). More specifically, at each round t, each agent i selects a decision $x_{i,t} \in \mathcal{X} \subseteq \mathbb{R}^{p}$ which only is a coordinate of the global decision. After all agents' selections, the convex loss function $f_t: \mathbb{R}^{np} \to \mathbb{R}$ is revealed to each agent and the agents suffer a loss $f_t(x_{1,t}, \ldots, x_{n,t})$. Same as the centralized scenario, the goal of the agents is to minimize the accumulated loss and the corresponding performance measure is regret which is defined with respect to the loss functions $\{f_t\}$. For this problem, Lee *et al.* (2018) achieved the optimal $\mathcal{O}(\sqrt{T})$ static regret bound.

Distributed online convex optimization with coupled constraints was considered in Yi *et al.* (2020) and Li *et al.* (2021). More specifically, at each round t, each agent i makes a decision $x_{i,t} \in \mathbb{X}_i \subseteq \mathbb{R}^{p_i}$. After making the selection, the convex local loss function $f_{i,t}: \mathbb{R}^{p_i} \to \mathbb{R}$ and constraint function $g_{i,t}: \mathbb{R}^{p_i} \to \mathbb{R}^{m}$ are revealed to agent i only, which respectively are a portion of the global loss function $f_t(x_t) = \sum_{i=1}^{n} f_{i,t}(x_{i,t})$ and coupled constraint function $g_t(x_t) = \sum_{i=1}^{n} g_i(x_{i,t})$, where $x_t = [x_{1,t}^T, \ldots, x_{n,t}^T]^T$ is the global decision. Same as the centralized scenario, the goal of the agents is to choose $\{x_{i,t}\}$ to minimize both regret and constraint violation, which are respectively defined with respect to the global loss functions $\{f_t\}$ and coupled constraint functions $\{g_t\}$. For this problem, Yi *et al.* (2020) achieved the same regret and constraint violation bounds as those achieved by the state-of-the-art centralized algorithms. When the constraint function is fixed, Li *et al.* (2021) showed that sublinear regret and constraint violation bounds can be achieved even though the communication graph is time-varying, directed and unbalanced.

Conclusions

This chapter reviewed online convex optimization which is a sequential decision making problem and introduced regret to measure the performance of online algorithms. It then presented variations with long-term and time-varying constraints, introducing an additional performance metric, constraint violation or the stricter cumulative constraint violation, to measure the violation of constraint functions. Moreover, these online convex optimization problems had also been extend to distributed settings to leverage the benefits of parallel computing. The differences among these scenarios were briefly discussed and state-of-the-art results in each scenario were summarized. While there has been significant progress in this direction, numerous open problems remain. For example, it is challenging to consider nonconvex loss functions since the standard metric, regret, is not suitable for online nonconvex optimization. Another interesting problem is reducing communication complexity for distributed online algorithms, with potential solutions including using periodic or event-triggered communication, as discussed in other chapters in the encyclopedia.

Acknowledgments

The work of Xinlei Yi and Karl H. Johansson was partially supported by the Knut and Alice Wallenberg Foundation through a postdoc grant and Wallenberg Scholar grant, and by the Swedish Research Council Distinguished Professor Grant 2017–01078. The work of Tao Yang was partially supported by the National Natural Science Foundation of China under Grants 62133003,61991403, 61991404 and 61991400.

References

Akbari, M., Gharesifard, B., Linder, T., 2017. Distributed online convex optimization on time-varying directed graphs. IEEE Transactions on Control of Network Systems 4 (3), 417–428.
Boyd, S., Vandenberghe, L., 2004. Convex Optimization. Cambridge University Press.
Cesa-Bianchi, N., Long, P.M., Warmuth, M.K., 1996. Worst-case quadratic loss bounds for prediction using linear functions and gradient descent. IEEE Transactions on Neural Networks 7 (3), 604–619.
Crammer, K., Dekel, O., Keshet, J., Shalev-Shwartz, S., Singer, Y., 2006. Online passive aggressive algorithms. Journal of Machine Learning Research 7, 551–585.
Goldfarb, A., Tucker, C., 2011. Online display advertising: Targeting and obtrusiveness. Marketing Science 30 (3), 389–404.
Guo, H., Liu, X., Wei, H., Ying, L., 2022. Online convex optimization with hard constraints: Towards the best of two worlds and beyond. Advances in Neural Information Processing Systems 36426–36439.
Hazan, E., 2016. Introduction to online convex optimization. Foundations and Trends in Optimization 2 (3-4), 157–325.
Hazan, E., Agarwal, A., Kale, S., 2007. Logarithmic regret algorithms for online convex optimization. Machine Learning 69 (2-3), 169–192.
Koloskova, A., Stich, S., Jaggi, M., 2019. Decentralized stochastic optimization and gossip algorithms with compressed communication. In: Proceedings of the International Conference on Machine Learning, pp. 3478–3487.
Lee, S., Nedić, A., Raginsky, M., 2018. Coordinate dual averaging for decentralized online optimization with nonseparable global objectives. IEEE Transactions on Control of Network Systems 5 (1), 34–44.
Li, X., Xie, L., Li, N., 2023. A survey on distributed online optimization and online games. Annual Reviews in Control 56, 100904.
Li, X., Yi, X., Xie, L., 2021. Distributed online optimization for multi-agent networks with coupled inequality constraints. IEEE Transactions on Automatic Control 66 (8), 3575–3591.
Mahdavi, M., Jin, R., Yang, T., 2012. Trading regret for efficiency: Online convex optimization with long term constraints. Journal of Machine Learning Research 13 (81), 2503–2528.
Mairal, J., Bach, F., Ponce, J., Sapiro, G., 2009. Online dictionary learning for sparse coding. In: Proceedings of the International Conference on Machine Learning, pp. 689–696.
Mateos-Núnez, D., Cortés, J., 2014. Distributed online convex optimization over jointly connected digraphs. IEEE Transactions on Network Science and Engineering 1 (1), 23–37.
Shalev-Shwartz, S., 2012. Online learning and online convex optimization. Foundations and Trends in Machine Learning 4 (2), 107–194.
Sun, W., Dey, D., Kapoor, A., 2017. Safety-aware algorithms for adversarial contextual bandit. In: Proceedings of the International Conference on Machine Learning, pp. 3280–3288.
Yi, X., Li, X., Xie, L., Johansson, K.H., 2020. Distributed online convex optimization with time-varying coupled inequality constraints. IEEE Transactions on Signal Processing 68, 731–746.
Yi, X., Li, X., Yang, T., *et al.*, 2021. Regret and cumulative constraint violation analysis for online convex optimization with long term constraints. In: Proceedings of the International Conference on Machine Learning, pp. 11998–12008.
Yi, X., Li, X., Yang, T., *et al.*, 2023a. Regret and cumulative constraint violation analysis for distributed online constrained convex optimization. IEEE Transactions on Automatic Control 68 (5), 2875–2890.
Yi, X., Li, X., Yang, T., *et al.*, 2023b. Distributed online convex optimization with adversarial constraints: Reduced cumulative constraint violation bounds under slater's condition. arXiv preprint arXiv:2306.00149.
Yu, H., Neely, M.J., 2020. A low complexity algorithm with $O(\sqrt{T})$ regret and $O(1)$ constraint violations for online convex optimization with long term constraints. Journal of Machine Learning Research 21 (1), 1–24.
Yu, H., Neely, M., Wei, X., 2017. Online convex optimization with stochastic constraints. Advances in Neural Information Processing Systems 1428–1438.
Yuan, D., Proutiere, A., Shi, G., 2022. Distributed online optimization with long-term constraints. IEEE Transactions on Automatic Control 67 (3), 1089–1104.
Yuan, J., Lamperski, A., 2018. Online convex optimization for cumulative constraints. Advances in Neural Information Processing Systems 6140–6149.
Zhang, L., Lu, S., Zhou, Z.H., 2018. Adaptive online learning in dynamic environments. Advances in Neural Information Processing Systems 1323–1333.
Zhao, P., Zhang, L., 2021. Improved analysis for dynamic regret of strongly convex and smooth functions. Learning for Dynamics and Control 48–59.
Zinkevich, M., 2003. Online convex programming and generalized infinitesimal gradient ascent. In: Proceedings of the International Conference on Machine Learning, pp. 928–936.