ROYAL INSTITUTE
OF TECHNOLOGY

Regret and Cumulative Constraint Violation Analysis for Online Convex Optimization with Long Term Constraints

Xinlei Yi
KTH Royal Institute of Technology, Sweden
xinleiy@kth.se

Joint work with Xiuxian Li, Tao Yang, Lihua Xie, Tianyou Chai, and Karl H. Johansson

Online Convex Optimization (OCO)

Online convex optimization [Zinkevich, ICML, 2003]

for $t = 1$ to T **do**

end for

Online Convex Optimization (OCO)

Online convex optimization [Zinkevich, ICML, 2003]

for $t = 1$ to T **do**

  Learner picks a decision $x_t \in \mathcal{X}$

  Adversary picks a convex function $f_t : \mathcal{X} \to \mathbb{R}$

end for

Online Convex Optimization (OCO)

Online convex optimization **[Zinkevich, ICML, 2003]**

1: **for** $t = 1$ to T **do**
2: Learner picks a decision $x_t \in \mathcal{X}$
Adversary picks a convex function $f_t : \mathcal{X} \to \mathbb{R}$
3: Learner suffers a loss $f_t(x_t)$ and updates x_t
4: **end for**

Cumulative loss (CL): $\sum_{t=1}^{T} f_t(x_t)$

Online Convex Optimization (OCO)

Online convex optimization [Zinkevich, ICML, 2003]

1: **for** $t = 1$ to T **do**
2: Learner picks a decision $x_t \in \mathcal{X}$
Adversary picks a convex function $f_t : \mathcal{X} \to \mathbb{R}$
3: Learner suffers a loss $f_t(x_t)$ and updates x_t
4: **end for**

Cumulative loss (CL): $\sum_{t=1}^{T} f_t(x_t)$

- Learner picks x_t **before** knowing f_t
- Learner is impossible to minimize CL directly

Online Convex Optimization (OCO)

Online convex optimization **[Zinkevich, ICML, 2003]**

1: **for** $t = 1$ to T **do**
2: Learner picks a decision $x_t \in \mathcal{X}$
Adversary picks a convex function $f_t : \mathcal{X} \to \mathbb{R}$
3: Learner suffers a loss $f_t(x_t)$ and updates x_t
4: **end for**

Cumulative loss (CL): $\sum_{t=1}^{T} f_t(x_t)$

- Learner picks x_t **before** knowing f_t
- Learner is impossible to minimize CL directly

Learner's objective: choose $\{x_t\}$ s.t. **regret** grows sublinearly

Online Convex Optimization (OCO)

Online convex optimization **[Zinkevich, ICML, 2003]**

1: **for** $t = 1$ to T **do**
2: Learner picks a decision $x_t \in \mathcal{X}$
Adversary picks a convex function $f_t : \mathcal{X} \to \mathbb{R}$
3: Learner suffers a loss $f_t(x_t)$ and updates x_t
4: **end for**

Learner's objective: choose $\{x_t\}$ s.t. **regret** grows sublinearly

$$\mathrm{Reg}(\{x_t\}, \{y_t\}) = \underbrace{\sum_{t=1}^{T} f_t(x_t)}_{\text{CL of an online learner}} - \underbrace{\sum_{t=1}^{T} f_t(y_t)}_{\text{CL of a comparator sequence}}$$

Online Convex Optimization (OCO)

Online convex optimization [Zinkevich, ICML, 2003]

1: **for** $t = 1$ to T **do**
2: Learner picks a decision $x_t \in \mathcal{X}$
Adversary picks a convex function $f_t : \mathcal{X} \to \mathbb{R}$
3: Learner suffers a loss $f_t(x_t)$ and updates x_t
4: **end for**

Learner's objective: choose $\{x_t\}$ s.t. **regret** grows sublinearly

$$\mathrm{Reg}(\{x_t\}, \{y_t\}) = \underbrace{\sum_{t=1}^{T} f_t(x_t)}_{\text{CL of an online learner}} - \underbrace{\sum_{t=1}^{T} f_t(y_t)}_{\text{CL of a comparator sequence}}$$

- **Static** regret: $\{y_t\} = \{x^*\}$, where $x^* = \arg\min_{x \in \mathcal{X}} \sum_{t=1}^{T} f_t(x)$
- **Dynamic** regret: $\{y_t\} = \{x_t^*\}$, where $x_t^* = \arg\min_{x_t \in \mathcal{X}} f_t(x_t)$

Optimal Regret Bound

Online gradient descent (OGD) [Zinkevich, ICML, 2003]

1: **for** $t = 1$ to T **do**
2: Learner picks a decision $x_t \in \mathcal{X}$
Adversary chooses a convex function $f_t : \mathcal{X} \to \mathbb{R}$
3: Learner suffers a loss $f_t(x_t)$
4: **end for**

Optimal Regret Bound

Online gradient descent (OGD) **[Zinkevich, ICML, 2003]**

1: **for** $t = 1$ to T **do**

2: Learner picks a decision $x_t \in \mathcal{X}$

Adversary chooses a convex function $f_t : \mathcal{X} \to \mathbb{R}$

3: Learner suffers a loss $f_t(x_t)$ and updates

$$x_{t+1} = \mathcal{P}_{\mathcal{X}}(x_t - \alpha \nabla f_t(x_t))$$

4: **end for**

- **Projection operator:** $\mathcal{P}_{\mathcal{X}}(x) = \arg\min_{y \in \mathcal{X}} \|x - y\|$

Optimal Regret Bound

Online gradient descent (OGD) [Zinkevich, ICML, 2003]

1: **for** $t = 1$ to T **do**

2: Learner picks a decision $x_t \in \mathcal{X}$

Adversary chooses a convex function $f_t : \mathcal{X} \to \mathbb{R}$

3: Learner suffers a loss $f_t(x_t)$ and updates

$$x_{t+1} = \mathcal{P}_{\mathcal{X}}(x_t - \alpha \nabla f_t(x_t))$$

4: **end for**

- **Projection operator:** $\mathcal{P}_{\mathcal{X}}(x) = \arg\min_{y \in \mathcal{X}} \|x - y\|$
- **Optimal static regret** bound: $\mathcal{O}(\sqrt{T})$

Optimal Regret Bound

Online gradient descent (OGD) [Zinkevich, ICML, 2003]

for $t = 1$ to T **do**
  Learner picks a decision $x_t \in \mathcal{X}$
  Adversary chooses a convex function $f_t : \mathcal{X} \to \mathbb{R}$
  Learner suffers a loss $f_t(x_t)$ and updates

$$x_{t+1} = \mathcal{P}_{\mathcal{X}}(x_t - \alpha \nabla f_t(x_t))$$

end for

- **Optimal static regret** bound: $\mathcal{O}(\sqrt{T})$

Ader [Zhang et al., NeurIPS, 2018]

- Run OGD multiple times in parallel, each with a different stepsize
- Choose the best one using an expert-tracking algorithm

Optimal Regret Bound

Online gradient descent (OGD) [Zinkevich, ICML, 2003]

1: **for** $t = 1$ to T **do**
2: Learner picks a decision $x_t \in \mathcal{X}$
Adversary chooses a convex function $f_t : \mathcal{X} \to \mathbb{R}$
3: Learner suffers a loss $f_t(x_t)$ and updates

$$x_{t+1} = \mathcal{P}_{\mathcal{X}}(x_t - \alpha \nabla f_t(x_t))$$

4: **end for**

- **Optimal static regret** bound: $\mathcal{O}(\sqrt{T})$

Ader [Zhang et al., NeurIPS, 2018]

- Run OGD multiple times in parallel, each with a different stepsize
- Choose the best one using an expert-tracking algorithm
- **Optimal regret** bound: $\mathcal{O}(\sqrt{T(1 + P_T)})$
- **Path-length** of comparator sequence: $P_T = \sum_{t=1}^{T-1} \|y_{t+1} - y_t\|$

OCO with Long Term Constraints

Online gradient descent: $x_{t+1} = \mathcal{P}_{\mathcal{X}}(x_t - \alpha \nabla f_t(x_t))$

Projection operator: $\mathcal{P}_{\mathcal{X}}(x) = \arg\min_{y \in \mathcal{X}} \|x - y\|$

Bottleneck is the computational cost of $\mathcal{P}_{\mathcal{X}}(\cdot)$

- $\mathcal{X}$ is a simple set, e.g., a box or a ball
- $\mathcal{X} = \{x : \ g(x) \leq \mathbf{0}_m, \ x \in \mathbb{X}\}$

OCO with Long Term Constraints

Bottleneck is the computational cost of $\mathcal{P}_{\mathcal{X}}(\cdot)$

- $\mathcal{X}$ is a simple set, e.g., a box or a ball
- $\mathcal{X} = \{x : \ g(x) \le \mathbf{0}_m, \ x \in \mathbb{X}\}$

$$\mathcal{X} = \{x : \underbrace{g(x) \le \mathbf{0}_m}_{\text{“Soft” constraint}}, \underbrace{x \in \mathbb{X}}_{\text{“Hard” constraint}}\}$$

- “Soft” constraint could be violated sometimes, but must be satisfied in the long run
- “Hard” constraint must be satisfied in each round

OCO with Long Term Constraints

Bottleneck is the computational cost of $\mathcal{P}_{\mathcal{X}}(\cdot)$

- $\mathcal{X}$ is a simple set, e.g., a box or a ball
- $\mathcal{X} = \{x : \ g(x) \leq \mathbf{0}_m, \ x \in \mathbb{X}\}$

$$\mathcal{X} = \{x : \underbrace{g(x) \leq \mathbf{0}_m}_{\text{"Soft" constraint}}, \underbrace{x \in \mathbb{X}}_{\text{"Hard" constraint}} \}$$

- "Soft" constraint could be violated sometimes, but must be satisfied in the long run
- "Hard" constraint must be satisfied in each round

OCO with long term constraints [Mahdavi et al., JMLR, 2012]

Learner's objective is to choose $\{x_t\}$ s.t. both **regret** and **constraint violation** grow sublinearly

- **Constraint violation:** $\|[\sum_{t=1}^{T} g(x_t)]_+\|$, where $[\cdot]_+ = \mathcal{P}_{\mathbb{R}^d_+}(\cdot)$

Stricter Metric: Cumulative Constraint Violation

Constraint violation: $\|[\sum_{t=1}^{T} g(x_t)]_+\|$

- **Drawback**: constraint violations at some rounds can be compensated by strictly feasible decisions at other rounds

Stricter Metric: Cumulative Constraint Violation

Constraint violation: $\|[\sum_{t=1}^{T} g(x_t)]_+\|$

- **Drawback**: constraint violations at some rounds can be compensated by strictly feasible decisions at other rounds

Cumulative constraint violation [Yuan & Lamperski, NeurIPS, 2018]:

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

State-of-the-Art

Reference	Loss functions	Static regret	Regret	Constraint violation	Cumulative constraint violation
Mahdavi et al. (2012)	Convex	$\mathcal{O}(\sqrt{T})$	Not given	$\mathcal{O}(T^{3/4})$	Not given
Jenatton et al. (2016)	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	Not given	$\mathcal{O}(T^{1-c/2})$	Not given
	Strongly convex	$\mathcal{O}(T^{c})$			
Yuan & Lamperski (2018)	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	Not given	$\mathcal{O}(T^{1-c/2})$	
	Strongly convex	$\mathcal{O}(\log(T))$		$\mathcal{O}(\sqrt{\log(T)T})$	
Yu & Neely (2020)	Convex	$\mathcal{O}(\sqrt{T})$	Not given	$\mathcal{O}(T^{1/4})$	Not given

Motivation

Reference	Loss functions	Static regret	Regret	Constraint violation	Cumulative constraint violation
Mahdavi et al. (2012)	Convex	$\mathcal{O}(\sqrt{T})$	**Not given**	$\mathcal{O}(T^{3/4})$	Not given
Jenatton et al. (2016)	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	**Not given**	$\mathcal{O}(T^{1-c/2})$	Not given
	Strongly convex	$\mathcal{O}(T^{c})$			
Yuan & Lamperski (2018)	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	**Not given**	$\mathcal{O}(T^{1-c/2})$	
	Strongly convex	$\mathcal{O}(\log(T))$		$\mathcal{O}(\sqrt{\log(T)T})$	
Yu & Neely (2020)	Convex	$\mathcal{O}(\sqrt{T})$	**Not given**	$\mathcal{O}(T^{1/4})$	Not given

Motivation

Reference	Loss functions	Static regret	Regret	Constraint violation	Cumulative constraint violation
Mahdavi et al. (2012)	Convex	$\mathcal{O}(\sqrt{T})$	**Not given**	$\mathcal{O}(T^{3/4})$	Not given
Jenatton et al. (2016)	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	**Not given**	$\mathcal{O}(T^{1-c/2})$	Not given
	Strongly convex	$\mathcal{O}(T^{c})$			
Yuan & Lamperski (2018)	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	**Not given**	$\mathcal{O}(T^{1-c/2})$	
	Strongly convex	$\mathcal{O}(\log(T))$		$\mathcal{O}(\sqrt{\log(T)T})$	
Yu & Neely (2020)	Convex	$\mathcal{O}(\sqrt{T})$	**Not given**	$\mathcal{O}(T^{1/4})$	Not given

Questions

(i) Can cumulative constraint violation be reduced?

(ii) Can optimal regret and sublinear cumulative constraint violation be achieved?

Basic Approach

Algorithm 1 [Yu &Neely, JMLR, 2020]

Input: $\alpha > 0$ and $\gamma > 0$.
Initialize: $q_0 = \mathbf{0}_m$ and $x_1 \in \mathbb{X}$.
for $t = 2, \dots$ **do**
Observe $\partial f_{t-1}(x_{t-1})$.
Update

$$q_{t-1} = \max\{-\gamma g(x_{t-1}), q_{t-2} + \gamma g(x_{t-1})\},$$
$$\hat{q}_{t-1} = q_{t-1} + \gamma g(x_{t-1}),$$
$$x_t = \arg\min_{x \in \mathbb{X}} \{\alpha \langle \partial f_{t-1}(x_{t-1}), x\rangle + \alpha \langle \hat{q}_{t-1}, \gamma g(x)\rangle + \|x - x_{t-1}\|^2\}.$$

end for
Output: $\{x_t\}$.

- $\text{Regret}(\{x_t\}, \{x^*\}) = \mathcal{O}(\sqrt{T})$
- $\|[\sum_{t=1}^T g(x_t)]_+\| = \mathcal{O}(T^{1/4})$

Basic Approach

Algorithm 1 [Yu &Neely, JMLR, 2020]

Input: $\alpha > 0$ and $\gamma > 0$.
Initialize: $q_0 = \mathbf{0}_m$ and $x_1 \in \mathbb{X}$.
for $t = 2, \dots$ **do**
Observe $\partial f_{t-1}(x_{t-1})$.
Update

$$q_{t-1} = \max\{-\gamma g(x_{t-1}), q_{t-2} + \gamma g(x_{t-1})\},$$
$$\hat{q}_{t-1} = q_{t-1} + \gamma g(x_{t-1}),$$
$$x_t = \arg\min_{x\in\mathbb{X}}\{\alpha\langle\partial f_{t-1}(x_{t-1}), x\rangle + \alpha\langle\hat{q}_{t-1}, \gamma g(x)\rangle + \|x - x_{t-1}\|^2\}.$$

end for
Output: $\{x_t\}$.

Algorithm 1

Input: $\{\alpha_t > 0\}$ and $\{\gamma_t > 0\}$.
Initialize: $q_0 = \mathbf{0}_m$ and $x_1 \in \mathbb{X}$.
for $t = 2, \dots$ **do**
Observe $\partial f_{t-1}(x_{t-1})$.
Update

$$q_{t-1} = q_{t-2} + [\gamma_t g(x_{t-1})]_+,$$
$$\hat{q}_{t-1} = q_{t-1} + [\gamma_t g(x_{t-1})]_+,$$
$$x_t = \arg\min_{x\in\mathbb{X}}\{\alpha_t\langle\partial f_{t-1}(x_{t-1}), x\rangle + \alpha_t\langle\hat{q}_{t-1}, \gamma_t[g(x)]_+\rangle + \|x - x_{t-1}\|^2\}.$$

end for
Output: $\{x_t\}$.

- $\text{Regret}(\{x_t\}, \{x^*\}) = \mathcal{O}(\sqrt{T})$
- $\|[\sum_{t=1}^T g(x_t)]_+\| = \mathcal{O}(T^{1/4})$

Properties of Basic Approach: Convex

Theorem 1

Let $\alpha_t = \alpha_0/t^c$ and $\gamma = \gamma_0/\sqrt{\alpha_t}$, where $\alpha_0 > 0$, $c \in (0,1)$, and $\gamma_0 \in (0, 1/(\sqrt{2}G)]$ are constants. Then,

$$\mathrm{Reg}(\{x_t\},\{y_t\}) = \mathcal{O}(\alpha_0 T^{1-c} + T^c(1+P_T)/\alpha_0),$$

$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(\sqrt{\alpha_0} T^{(1-c)/2}).$$

Properties of Basic Approach: Convex

Theorem 1

Let $\alpha_t = \alpha_0/t^c$ and $\gamma = \gamma_0/\sqrt{\alpha_t}$, where $\alpha_0 > 0$, $c \in (0,1)$, and $\gamma_0 \in (0, 1/(\sqrt{2}G)]$ are constants. Then,

$$\mathrm{Reg}(\{x_t\},\{y_t\}) = \mathcal{O}(\alpha_0 T^{1-c} + T^c(1+P_T)/\alpha_0),$$
$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(\sqrt{\alpha_0} T^{(1-c)/2}).$$

(i) Choosing $\alpha_0 = 1$ yields

$$\mathrm{Reg}(\{x_t\},\{x^*\}) = \mathcal{O}(T^{\max\{c,1-c\}}),$$
$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(T^{(1-c)/2}) < \mathcal{O}(T^{1-c/2}).$$

State-of-the-art result [Yuan & Lamperski, NeurIPS, 2018]

$$\mathrm{Reg}(\{x_t\},\{x^*\}) = \mathcal{O}(T^{\max\{c,1-c\}}),$$
$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(T^{1-c/2}).$$

Properties of Basic Approach: Convex

Theorem 1

Let $\alpha_t = \alpha_0/t^c$ and $\gamma = \gamma_0/\sqrt{\alpha_t}$, where $\alpha_0 > 0$, $c \in (0,1)$, and $\gamma_0 \in (0, 1/(\sqrt{2}G)]$ are constants. Then,

$$\mathrm{Reg}(\{x_t\},\{y_t\}) = \mathcal{O}(\alpha_0 T^{1-c} + T^c(1+P_T)/\alpha_0),$$
$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(\sqrt{\alpha_0} T^{(1-c)/2}).$$

(i) Choosing $\alpha_0 = 1$ and $c = 0.5$ yields

$$\mathrm{Reg}(\{x_t\},\{x^*\}) = \mathcal{O}(\sqrt{T}),$$ Optimal static regret bound

$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(T^{1/4}).$$ Cumulative constraint violation

State-of-the-art result [Yu & Neely, JMLR, 2020]

$$\mathrm{Reg}(\{x_t\},\{x^*\}) = \mathcal{O}(\sqrt{T}),$$
$$\|[\sum_{t=1}^{T} g(x_t)]_+\| = \mathcal{O}(T^{1/4}).$$ Constraint violation

Properties of Basic Approach: Convex

Theorem 1

Let $\alpha_t = \alpha_0/t^c$ and $\gamma = \gamma_0/\sqrt{\alpha_t}$, where $\alpha_0 > 0$, $c \in (0,1)$, and $\gamma_0 \in (0, 1/(\sqrt{2}G)]$ are constants. Then,

$$\mathrm{Reg}(\{x_t\},\{y_t\}) = \mathcal{O}(\alpha_0 T^{1-c} + T^c(1+P_T)/\alpha_0),$$

$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(\sqrt{\alpha_0} T^{(1-c)/2}).$$

(ii) If P_T is known, choosing $\alpha_0 = \sqrt{1+P_T}$ and $c = 0.5$ yields

$$\mathrm{Reg}(\{x_t\},\{y_t\}) = \mathcal{O}(\sqrt{T(1+P_T)}),$$ Optimal regret bound

$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(T^{1/4}(1+P_T)^{1/4}) \le \mathcal{O}(\sqrt{T}).$$

Properties of Basic Approach: Strongly Convex

Recall: Theorem 1

Let $\alpha_t = \alpha_0/\sqrt{t}$ and $\gamma = \gamma_0/\sqrt{\alpha_t}$, where $\alpha_0 > 0$ and $\gamma_0 \in (0, 1/(\sqrt{2}G)]$ are constants. Then,

$$\mathrm{Reg}(\{x_t\}, \{x^*\}) = \mathcal{O}(\sqrt{T}),\ \sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(T^{1/4}).$$

Theorem 2

Suppose each $f_t(\cdot)$ is strongly convex. Let $\alpha_t = 1/(\mu t)$ and $\gamma = \gamma_0/\sqrt{\alpha_t}$, where $\gamma_0 \in (0, 1/(\sqrt{2}G)]$ are constants. Then,

$$\mathrm{Reg}(\{x_t\}, \{x^*\}) = \mathcal{O}(\log(T)),\ \sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(\log(T)).$$

State-of-the-art result [Yuan & Lamperski, NeurIPS, 2018]

Each $f_t(\cdot)$ is strongly convex:

$$\mathrm{Reg}(\{x_t\}, \{x^*\}) = \mathcal{O}(\log(T)),\ \sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(\sqrt{T\log(T)}).$$

Improved Approach: Motivation

Recall: Theorem 1

If P_T is known in advance, choosing $\alpha_t = \sqrt{(1+P_T)/t}$ yields

$$\mathrm{Reg}(\{x_t\},\{y_t\}) = \mathcal{O}(\sqrt{T(1+P_T)}),$$
$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(T^{1/4}(1+P_T)^{1/4}).$$

Improved Approach: Motivation

Recall: Theorem 1

If P_T is known in advance, choosing $\alpha_t = \sqrt{(1+P_T)/t}$ yields

$$\mathrm{Reg}(\{x_t\},\{y_t\}) = \mathcal{O}(\sqrt{T(1+P_T)}),$$

$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(T^{1/4}(1+P_T)^{1/4}).$$

- Question: $P_T = \sum_{t=1}^{T-1} \|y_{t+1} - y_t\|$ is normally unknown

Improved Approach: Motivation

Recall: Theorem 1

If P_T is known in advance, choosing $\alpha_t = \sqrt{(1+P_T)/t}$ yields

$$\mathrm{Reg}(\{x_t\},\{y_t\}) = \mathcal{O}(\sqrt{T(1+P_T)}),$$

$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(T^{1/4}(1+P_T)^{1/4}).$$

- Question: $P_T = \sum_{t=1}^{T-1} \|y_{t+1} - y_t\|$ is normally unknown

Solution

- Design $N = \log_2(\sqrt{T})$ different stepsizes $\alpha_{i,t} = 2^{i-1}/\sqrt{t}$:
  $\exists i_0 \leq N$, s.t. $\alpha_{i_0,t}$ is close to the optimal stepsize $\sqrt{(1+P_T)/t}$
- Run Algorithm 1 N times in parallel, each with stepsize $\alpha_{i,t}$
- Choose the optimal one using an expert-tracking algorithm

Improved Approach

Algorithm 2

Input: $N \in \mathbb{N}_+$, $\beta > 0$, $\{\alpha_{i,t} > 0\}$ and $\{\gamma_{i,t} > 0\}$.
Initialize: $q_{i,0} = \mathbf{0}_m$, $x_{i,1} \in \mathbb{X}$, $w_{i,1} = \frac{N+1}{i(i+1)N}$, $\forall i \in [N]$, and $x_1 = \sum_{i=1}^{N} w_{i,1} x_{i,1}$.
for $t = 2, \dots$ **do**
Observe $\partial f_{t-1}(x_{t-1})$.
Update

$$\begin{aligned} q_{i,t-1} &= q_{i,t-2} + [\gamma_{i,t} g(x_{i,t-1})]_+, \\ \hat{q}_{i,t-1} &= q_{i,t-1} + [\gamma_{i,t} g(x_{i,t-1})]_+, \\ x_{i,t} &= \arg\min_{x \in \mathbb{X}} \{\alpha_{i,t} \langle \partial f_{t-1}(x_{t-1}), x \rangle \\ &\quad + \alpha_{i,t} \langle \hat{q}_{i,t-1}, [\gamma_{i,t} g(x)]_+ \rangle + \|x - x_{i,t-1}\|^2\}, \end{aligned}$$

Algorithm 1

$$\begin{aligned} l_{i,t-1} &= \langle \partial f_{t-1}(x_{t-1}), x_{i,t-1} - x_{t-1} \rangle, \\ w_{i,t} &= \frac{w_{i,t-1} e^{-\beta l_{i,t-1}}}{\sum_{i=1}^{N} w_{i,t-1} e^{-\beta l_{i,t-1}}}, \\ x_t &= \sum_{i=1}^{N} w_{i,t} x_{i,t}. \end{aligned}$$

Expert-tracking

end for
Output: $\{x_t\}$.

Properties of Improved Approach

Theorem 3 (Optimal Regret)

Let $N = \log_2(\sqrt{1+T})$, $\beta = 1/\sqrt{T}$, $\alpha_{i,t} = 2^{i-1}/\sqrt{t}$, and $\gamma = \gamma_0/\sqrt{\alpha_t}$, where $\gamma_0 \in (0, 1/(\sqrt{2}G)]$ is a constant. Then,

$$\mathrm{Reg}(\{x_t\}, \{y_t\}) = \mathcal{O}(\sqrt{T(1+P_T)}),$$ Optimal regret bound

$$\sum_{t=1}^{T} \|[g(x_t)]_+\| = \mathcal{O}(\sqrt{T}).$$

Online Linear Programming

- $\mathbb{X} \subseteq \mathbb{R}^p$, $f_t(x) = \langle \theta_t, x \rangle$ and $g(x) = Ax - b$
- The settings on p, $\mathbb{X}$, θ_t, A, and b are similar to Yu & Neely (2020)

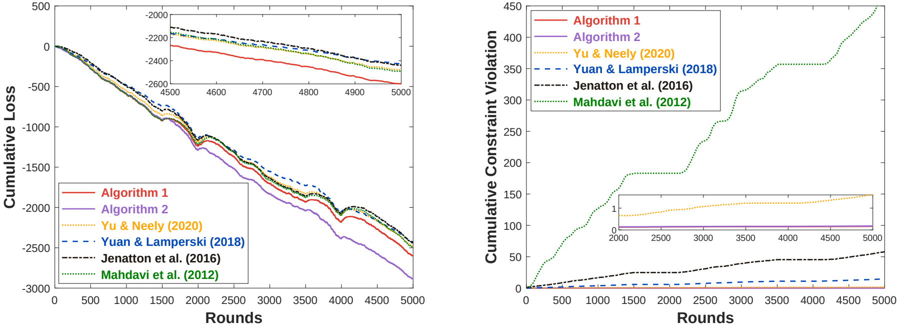

Online Quadratic Programming

- $\mathbb{X} \subseteq \mathbb{R}^p$, $f_t(x) = \|x - \theta_t\|^2 + 20\langle \theta_t, x \rangle$ and $g(x) = Ax - b$

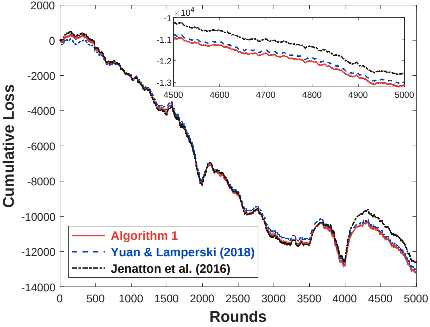

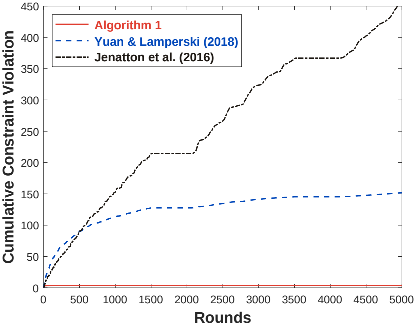

Conclusions

Reference	Loss functions	Static regret	Regret	Constraint violation	Cumulative constraint violation
Mahdavi et al. (2012)	Convex	$\mathcal{O}(\sqrt{T})$	**Not given**	$\mathcal{O}(T^{3/4})$	Not given
Jenatton et al. (2016)	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	**Not given**	$\mathcal{O}(T^{1-c/2})$	Not given
	Strongly convex	$\mathcal{O}(T^{c})$			
Yuan & Lamperski (2018)	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	**Not given**	$\mathcal{O}(T^{1-c/2})$	
	Strongly convex	$\mathcal{O}(\log(T))$		$\mathcal{O}(\sqrt{\log(T)T})$	
Yu & Neely (2020)	Convex	$\mathcal{O}(\sqrt{T})$	**Not given**	$\mathcal{O}(T^{1/4})$	Not given
Algorithm 1	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	$\mathcal{O}(\sqrt{T}(1+P_T))$	$\mathcal{O}(T^{(1-c)/2})$	
	Strongly convex	$\mathcal{O}(\log(T))$	Not given	$\mathcal{O}(\log(T))$	
Algorithm 2	Convex	$\mathcal{O}(T^{\max\{c,1-c\}})$	$\mathcal{O}(\sqrt{T(1+P_T)})$	$\mathcal{O}(\sqrt{T})$	

Future work

- Reduce regret under strong convexity and/or smoothness condition
- Reduce cumulative constraint violation under the Slater condition

Thanks for your time!

Jenatton, Rodolphe, Jim Huang, and Cédric Archambeau (2016). "Adaptive algorithms for online convex optimization with long-term constraints". In: *International Conference on Machine Learning*, pp. 402–411.

Mahdavi, Mehrdad, Rong Jin, and Tianbao Yang (2012). "Trading regret for efficiency: Online convex optimization with long term constraints". In: *Journal of Machine Learning Research* 13.81, pp. 2503–2528.

Yu, Hao and Michael J. Neely (2020). "A Low Complexity Algorithm with $O(\sqrt{T})$ Regret and $O(1)$ Constraint Violations for Online Convex Optimization with Long Term Constraints". In: *Journal of Machine Learning Research* 21.1, pp. 1–24.

Yuan, Jianjun and Andrew Lamperski (2018). "Online convex optimization for cumulative constraints". In: *Advances in Neural Information Processing Systems*, pp. 6140–6149.

Zhang, Lijun, Shiyin Lu, and Zhi-Hua Zhou (2018). "Adaptive online learning in dynamic environments". In: *Advances in Neural Information Processing Systems*, pp. 1323–1333.

Zinkevich, Martin (2003). "Online convex programming and generalized infinitesimal gradient ascent". In: *International Conference on Machine Learning*, pp. 928–936.