The 39th Annual Conference on Learning Theory

A Simple, Optimal and Efficient Algorithm for Online Exp-Concave Optimization



Yi-Han Wang   Peng Zhao   Zhi-Hua Zhou

LAMDA Group

Nanjing University, China

Outline

- **Problem Setting**
  - Online Exp-Concave Optimization
  - Online Newton Step (ONS)
  - A COLT'13 Open Problem
- **LightONS**
  - Techniques
  - Answering COLT'13 Open Problem
  - Applications to Downstream Results
- **Conclusion**

Online Convex Optimization (OCO)

- **Online Learning**: *interactive game* between *learner* and *environment*

> Each round $t = 1, 2, \ldots, T$:
> - *learner* picks a decision $\mathbf{x}_t \in \mathcal{X} \subseteq \mathbb{R}^d$
> - at the same time, ***environment*** picks convex loss $f_t : \mathcal{X} \to \mathbb{R}$
> - the *learner* suffers $f_t(\mathbf{x}_t)$ and observes $\nabla f_t(\mathbf{x}_t)$

- **Regret**: *learner*'s excess loss compared to ***best offline model in hindsight***

$$\mathrm{Reg}_T \triangleq \boxed{\sum_{t=1}^{T} f_t(\mathbf{x}_t)} - \boxed{\min_{\mathbf{u}\in\mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{u})}$$

Online eXp-concave Optimization (OXO)

Definition (exponential concavity). $f : \mathcal{X} \to \mathbb{R}$ is α-exp-concave ($\alpha > 0$) if and only if $\exp(-\alpha f(\cdot))$ is concave on $\mathcal{X}$. [Kivinen, Warmuth, COLT'99]

- **Many Applications**: linear / logistic regression; portfolio selection; LQR control; ...
- **Minimax Regret**:
  - OCO: $\mathrm{Reg}_T = \Theta(\sqrt{T})$
  - OXO: $\mathrm{Reg}_T = \Theta(d \log T)$
- **Rich Structures**: *local geometry of losses* (e.g., *local norms*)

Online Newton Step (ONS)

[Hazan, Kalai, Kale, Agarwal, COLT'06]

ONS

Initialize $\mathbf{x}_1 \in \mathcal{X}$, $A_0 = \epsilon I$.

Each round $t = 1, 2, \ldots, T$:

- Play $\mathbf{x}_t$, observe $\nabla f_t(\mathbf{x}_t)$
- $A_t = A_{t-1} + \nabla f_t(\mathbf{x}_t) \nabla f_t(\mathbf{x}_t)^\top$
- $\widehat{\mathbf{x}}_{t+1} = \mathbf{x}_t - \frac{1}{\gamma_0} A_t^{-1} \nabla f_t(\mathbf{x}_t)$
- $\mathbf{x}_{t+1} = \Pi_{\mathcal{X}}^{A_t}\left[\widehat{\mathbf{x}}_{t+1}\right]$

Newton-style update with Hessian-like matrix

$$A_t = \epsilon I + \sum_{s=1}^{t} \nabla f_s(\mathbf{x}_s) \nabla f_s(\mathbf{x}_s)^\top$$

Mahalanobis projection

$$\Pi_{\mathcal{X}}^{A}[\mathbf{y}] \triangleq \min_{\mathbf{x} \in \mathcal{X}} (\mathbf{x} - \mathbf{y})^\top A (\mathbf{x} - \mathbf{y})$$

Lemma (local norm). If f is α-exp-concave and G-Lipschitz on $\mathcal{X}$, then $\forall\, \mathbf{x}, \mathbf{u} \in \mathcal{X}$:

$$f(\mathbf{x}) - f(\mathbf{u}) \le \nabla f(\mathbf{x})^\top (\mathbf{x} - \mathbf{u}) - \frac{\gamma_0}{2} \left(\nabla f(\mathbf{x})^\top (\mathbf{x} - \mathbf{u})\right)^2,$$

where $\gamma_0 = \frac{1}{2} \min\{\frac{1}{DG}, \alpha\}$ and $D \ge \max_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|_2$.

$(\mathbf{x}_t - \mathbf{u})^\top \nabla f_t(\mathbf{x}_t) \nabla f_t(\mathbf{x}_t)^\top (\mathbf{x}_t - \mathbf{u})$

Online Newton Step (ONS) [Hazan, Kalai, Kale, Agarwal, COLT'06]

ONS

Initialize $\mathbf{x}_1 \in \mathcal{X}$, $A_0 = \epsilon I$.

Each round $t = 1, 2, \ldots, T$:

- Play $\mathbf{x}_t$, observe $\nabla f_t(\mathbf{x}_t)$
- $A_t = A_{t-1} + \nabla f_t(\mathbf{x}_t) \nabla f_t(\mathbf{x}_t)^\mathsf{T}$
- $\widehat{\mathbf{x}}_{t+1} = \mathbf{x}_t - \frac{1}{\gamma_0} A_t^{-1} \nabla f_t(\mathbf{x}_t)$
- $\mathbf{x}_{t+1} = \boxed{\Pi_{\mathcal{X}}^{A_t}\left[\widehat{\mathbf{x}}_{t+1}\right]}$

- **Computational Challenge**:

  even for very simple domains, including *the unit ball*, each ***Mahalanobis projection*** costs $\widetilde{O}(d^3)$ arithmetic operations.

  $\Longrightarrow$ ***total runtime*** $\widetilde{O}(d^3 T)$

* $\widetilde{O}(d^{2.3714})$ when exploiting fast matrix multiplication.

Theorem (ONS regret). ONS achieves *minimax optimal regret*:

$$\mathrm{Reg}_T \leq \frac{d}{2\gamma_0} \log\left(1 + \frac{G^2 T}{d\epsilon}\right) + \frac{\gamma_0 \epsilon D^2}{8} = O(d \log T).$$

A COLT'13 Open Problem [Koren, COLT'13]

Consider **Stochastic eXp-concave Optimization** (**SXO**).

$$\text{Objective:} \quad \text{ExcessRisk} \leq \varepsilon, \quad \text{where}$$

$$\text{ExcessRisk} \triangleq \mathbb{E}_f[f(\bar{\mathbf{x}}_T)] - \min_{\mathbf{u} \in \mathcal{X}} \mathbb{E}_f[f(\mathbf{u})]$$

- With ***online-to-batch conversion (O2BC)***, ONS incurs runtime $\tilde{O}(d^4/\varepsilon)$.

Is it possible to perform any better?

JMLR: Workshop and Conference Proceedings vol 30 (2013) 1–3

Open Problem: Fast Stochastic Exp-Concave Optimization

Tomer Koren TOMERK@TECHNION.AC.IL
Technion—Israel Institute of Technology

Abstract

Stochastic exp-concave optimization is an important primitive in machine learning that captures several fundamental problems, including linear regression, logistic regression and more. The exp-concavity property allows for fast convergence rates, as compared to general stochastic optimization. However, current algorithms that attain such rates scale poorly with the dimension n and run in time $O(n^4)$, even on very simple instances of the problem. The question we pose is whether it is possible to obtain fast rates for exp-concave functions using more computationally-efficient algorithms.

Open Problem: Is it possible to find an optimization algorithm that attains the rate of $\tilde{O}(n/\varepsilon)$ for exp-concave objectives, with only linear-time computation per iteration? Is it possible to perform any better than $\tilde{O}(n^4/\varepsilon)$ overall?

Partial Answer: OQNS [Mhammedi and Gatmiry, COLT'23]

Key Steps of Online Quasi-Newton Steps (OQNS)

circumvent ***Mahalanobis projections*** with ***log barrier***

$\mathbf{x}_{t+1} = \mathbf{x}_t - \text{Approximate}\big((\nabla^2 \Phi_t(\mathbf{x}_t))^{-1} \nabla \Phi_t(\mathbf{x}_t)\big)$ where

$$\Phi_t(\mathbf{x}) \triangleq \eta d \log \frac{1}{1-\|\mathbf{x}\|^2} + \frac{d+\eta}{2}\|\mathbf{x}\|^2 + \sum_{s=1}^{t}\left(\nabla f_s(\mathbf{x}_s)^\top \mathbf{x} + \frac{\gamma}{2}\big(\nabla f_s(\mathbf{x}_s)^\top(\mathbf{x}-\mathbf{x}_s)\big)^2\right).$$

- OQNS achieves ***minimax optimal regret*** $O(d \log T)$
- OQNS reduces runtime to $O(d^2 T \log T + d^3 \sqrt{T \log T})$ for the unit ball

Critical Issue: OQNS is ***not*** "***ONS-like***" (does ***not*** apply to ***downstream results*** where ONS is backbone)!

Our Result

Algorithm	Regret	Runtime	ONS-Like
OGD	$\sqrt{T}$	dT	N/A
ONS	$d \log T$	$d^2 T + d^3 T \log T$	**Yes**
OQNS	$d \log T$ (constants $\times 25$)	$d^2 T \log T + d^3 \sqrt{T \log T}$	**No**
LightONS	$d \log T$ (same constants as ONS)	$d^2 T + d^3 \sqrt{T \log T}$	**Yes**

LightONS enjoys all favorable properties:
- ① *optimal regret*
- ② *reduced runtime*
- ③ *"ONS-like"*

LightONS.Core

LightONS.Core

Initialize $\mathbf{x}_1 \in \widetilde{\mathcal{X}}_k$, $A_0 = \epsilon I$.

Each round $t = 1, 2, \ldots, T$:

- Play $\mathbf{x}_t$, observe $\nabla f_t(\mathbf{x}_t)$
- $A_t = A_{t-1} + \nabla f_t(\mathbf{x}_t) \nabla f_t(\mathbf{x}_t)^\mathsf{T}$
- $\widehat{\mathbf{x}}_{t+1} = \mathbf{x}_t - \frac{1}{\gamma_0} A_t^{-1} \nabla f_t(\mathbf{x}_t)$
- $\mathbf{x}_{t+1} = \begin{cases} \widehat{\mathbf{x}}_{t+1} & \text{if } \|\widehat{\mathbf{x}}_{t+1}\|_2 \leq kD/2 \\ \Pi_{\widetilde{\mathcal{X}}_k}^{A_t}\left[\widehat{\mathbf{x}}_{t+1}\right] & \text{otherwise} \end{cases}.$

only 1 line **different from ONS**:

- ***deferred projection***

reduces projection count to $O(\sqrt{T})$

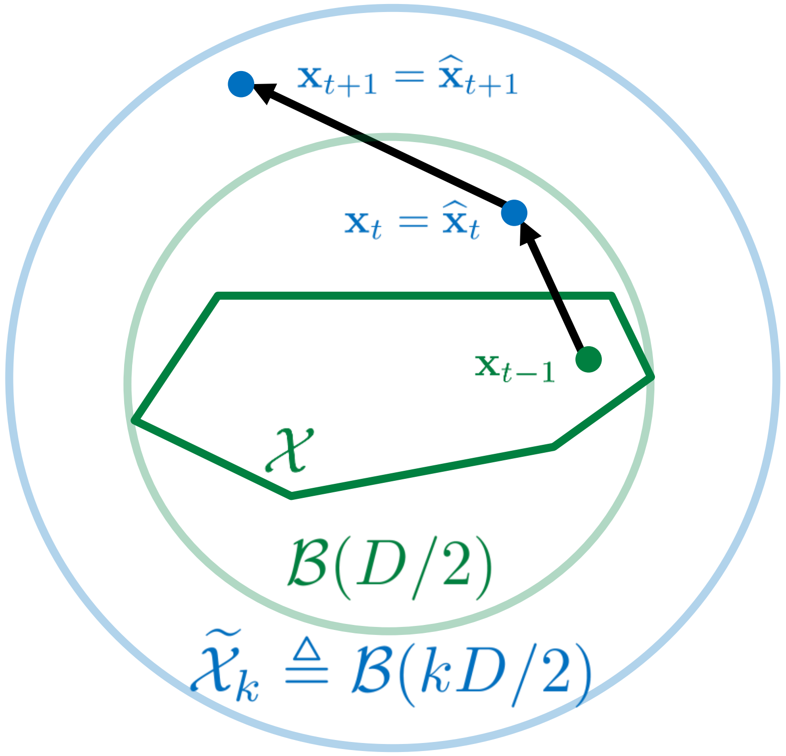

Improper Learning Issue

Theorem (LightONS.Core guarantees). With $k = 2$, LightONS.Core achieves regret $O(d \log T)$ with at most $O(\sqrt{T})$ Mahalanobis projections.

However, LightONS.Core regret is ***improper*** (“***unfair***”):

$$\mathrm{Reg}_T \triangleq \boxed{\sum_{t=1}^{T} f_t(\mathbf{x}_t)} - \boxed{\min_{\mathbf{u} \in \mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{u})}$$

$\mathbf{x}_t \in \widetilde{\mathcal{X}}_k \supset \mathcal{X}$ $\qquad$ $\mathbf{u} \in \mathcal{X}$

Critical Issue: improper learning ***may reduce the intrinsic difficulty*** of OXO.

Example: for logistic regression, proper regret $\Omega(de^G \log T)$, improper regret $O(d \log(GT))$.

[Hazan, Koren, Levy, COLT'14; Foster, Kale, Luo, Mohri, Sridharan, COLT'18]

LightONS

LightONS

Initialize $\mathbf{y}_1 \in \widetilde{\mathcal{X}}_k$, $A_0 = \epsilon I$.

Each round $t = 1, 2, \ldots, T$:

- Play $\mathbf{x}_t = \Pi_{\mathcal{X}}[\mathbf{y}_t]$, observe $\nabla f_t(\mathbf{x}_t)$, compute $\nabla g_t(\mathbf{y}_t)$
- $A_t = A_{t-1} + \nabla g_t(\mathbf{y}_t) \nabla g_t(\mathbf{y}_t)^\top$
- $\widehat{\mathbf{y}}_{t+1} = \mathbf{y}_t - \frac{1}{\gamma_0} A_t^{-1} \nabla g_t(\mathbf{y}_t)$
- $\mathbf{y}_{t+1} = \begin{cases} \widehat{\mathbf{y}}_{t+1} & \text{if } \|\widehat{\mathbf{y}}_{t+1}\|_2 \le kD/2 \\ \Pi_{\widetilde{\mathcal{X}}_k}^{A_t}[\widehat{\mathbf{y}}_{t+1}] & \text{otherwise} \end{cases}$.

only 2 lines **different from ONS**:

- ***deferred projection***

  reduces projection count to $O(\sqrt{T})$

- ***domain conversion***

  makes LightONS proper learning

Domain Conversion [Cutkosky and Orabona, COLT'18; Cutkosky, ICML'20]

- Run LightONS.Core on ***surrogate loss*** from parameter-free online learning

$$g_t(\mathbf{y}) = \nabla f_t(\mathbf{x}_t)^\top \mathbf{y} + \frac{\max\left\{-\nabla f_t(\mathbf{x}_t)^\top(\mathbf{y}_t - \mathbf{x}_t), 0\right\}}{\|\mathbf{y}_t - \mathbf{x}_t\|_2} \|\mathbf{y} - \Pi_{\mathcal{X}}[\mathbf{y}]\|_2$$

- ***Euclidean projection*** gives ***proper decisions***: $\mathbf{x}_t = \Pi_{\mathcal{X}}[\mathbf{y}_t]$

Computationally Cheap: only $O(d)$ extra time per round

Euclidean projection is much cheaper than Mahalanobis projections

Caveat: g_t is ***not*** exp-concave!

Domain Conversion [Cutkosky and Orabona, COLT'18; Cutkosky, ICML'20]

Fortunately, we prove that g_t inherits the ***local norm*** structure of f_t

Thus, LightONS shares the analysis with ONS

Master Lemma. If f is α-exp-concave and G-Lipschitz on $\mathcal{X}$, then $\forall\, \mathbf{x}, \mathbf{u} \in \mathcal{X}$:

$$f(\mathbf{x}) - f(\mathbf{u}) \leq \nabla f(\mathbf{x})^\top (\mathbf{x} - \mathbf{u}) - \frac{\gamma_0}{2} \left(\nabla f(\mathbf{x})^\top (\mathbf{x} - \mathbf{u})\right)^2$$
$$\leq \nabla g(\mathbf{y})^\top (\mathbf{y} - \mathbf{u}) - \frac{\gamma_0}{2} \left(\nabla g(\mathbf{y})^\top (\mathbf{y} - \mathbf{u})\right)^2,$$

where $\gamma_0 = \frac{1}{2} \min\{\frac{1}{DG}, \alpha\}$ and $D \geq \max_{\mathbf{x},\mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|_2$. $(\mathbf{y}_t - \mathbf{u})^\top \nabla g_t(\mathbf{y}_t) \nabla g_t(\mathbf{y}_t)^\top (\mathbf{y}_t - \mathbf{u})$

Answering COLT'13 Open Problem [Koren, COLT'13]

- Consider **Stochastic eXp-concave Optimization** (**SXO**).

$$\text{Objective:}\quad \text{ExcessRisk} \le \varepsilon,\quad \text{where}$$

$$\text{ExcessRisk} \triangleq \mathop{\mathbb{E}}_{f}[f(\bar{\mathbf{x}}_T)] - \min_{\mathbf{u}\in\mathcal{X}} \mathop{\mathbb{E}}_{f}[f(\mathbf{u})]$$

- With ***online-to-batch conversion (O2BC)***, ONS incurs runtime $\tilde{O}(d^4/\varepsilon)$.

Is it possible to perform any better?

JMLR: Workshop and Conference Proceedings vol 30 (2013) 1–3

Open Problem: Fast Stochastic Exp-Concave Optimization

Tomer Koren TOMERK@TECHNION.AC.IL
Technion—Israel Institute of Technology

Abstract

Stochastic exp-concave optimization is an important primitive in machine learning that captures several fundamental problems, including linear regression, logistic regression and more. The exp-concavity property allows for fast convergence rates, as compared to general stochastic optimization. However, current algorithms that attain such rates scale poorly with the dimension n and run in time $O(n^4)$, even on very simple instances of the problem. The question we pose is whether it is possible to obtain fast rates for exp-concave functions using more computationally-efficient algorithms.

Open Problem: Is it possible to find an optimization algorithm that attains the rate of $\tilde{O}(n/\varepsilon)$ for exp-concave objectives, with only linear-time computation per iteration? Is it possible to perform any better than $\tilde{O}(n^4/\varepsilon)$ overall?

- With ***O2BC***, LightONS achieves $\tilde{O}(d^3/\varepsilon)$ runtime ***in expect.*** and ***with high prob.***.

LightONS

LightONS

Initialize $\mathbf{y}_1 \in \widetilde{\mathcal{X}}_k$, $A_0 = \epsilon I$.

Each round $t = 1, 2, \ldots, T$:

- Play $\mathbf{x}_t = \Pi_{\mathcal{X}}[\mathbf{y}_t]$, observe $\nabla f_t(\mathbf{x}_t)$, compute $\nabla g_t(\mathbf{y}_t)$
- $A_t = A_{t-1} + \nabla g_t(\mathbf{y}_t) \nabla g_t(\mathbf{y}_t)^\top$
- $\widehat{\mathbf{y}}_{t+1} = \mathbf{y}_t - \frac{1}{\gamma_0} A_t^{-1} \nabla g_t(\mathbf{y}_t)$
- $\mathbf{y}_{t+1} = \begin{cases} \widehat{\mathbf{y}}_{t+1} & \text{if } \|\widehat{\mathbf{y}}_{t+1}\|_2 \leq kD/2 \\ \Pi_{\widetilde{\mathcal{X}}_k}^{A_t}[\widehat{\mathbf{y}}_{t+1}] & \text{otherwise} \end{cases}.$

only 2 lines **different from ONS**:

- ***deferred projection***

  reduces projection count to $O(\sqrt{T})$

- ***domain conversion***

  makes LightONS proper learning

- Both are "***orthogonal***" to ONS

  LightONS is "***ONS-like***"

LightONS is "ONS-Like"

① **Gradient-Norm Adaptivity**: ***problem-dependent regret*** sacles with $G_T \triangleq \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t)\|_2^2 \leq G^2 T$, instead of T. For ***unbounded online convex optimization***, this implies comparator-norm adaptivity (a.k.a. "parameter-free"), i.e., $\mathrm{Reg}_T \leq \widetilde{O}(\|\mathbf{u}\|_2 \sqrt{d G_T})$. [Cutkosky and Orabona, COLT'18]

② **Jointly Efficient Generalized Linear Bandits**: e.g., logistic bandits, K-arm, $0/1$-feedback, $\mathbb{P}(y_t = 0 \mid \mathbf{x}_t) = 1/(1+\exp(\mathbf{w}^\top \mathbf{x}_t))$. Challenge: simultaneously achieve *(a)* nice regret dependence on the condition number $\kappa \propto \exp(D)$, *(b)* one pass. [Zhang, Xu, Zhao, Sugiyama, NeurIPS'25]

③ **Memory-Efficient OXO**: only store top d' eigenvectors of the Hessian-like matrix A_t in ONS ($d' \ll d$), reduce memory from $O(d^2)$ to $O(d'd)$. [Luo, Agarwal, Cesa-Bianchi, Langford, NIPS'16]

	Regret (ONS & LightONS)	**Runtime (ONS)**	**Runtime (LightONS)**	**OQNS?**
①	$O(d \log G_T)$	$\widetilde{O}(d^3 T)$	$O(d^2 T + d^3 \sqrt{T} \log T)$	**Hardly**
②	$\widetilde{O}(d\sqrt{T} + \kappa d^2)$	$\widetilde{O}(d^2 K T + d^3 T)$	$\widetilde{O}(d^2 K T + d^3 \cdot \min\{\sqrt{\kappa d T} \log \kappa, T\})$	**Hardly**
③	$O(d' \log T)$	$\widetilde{O}(d^3 T)$	$\widetilde{O}(d' d T + d^3 \sqrt{T})$	**Intricate**

Conclusion

- **For Online eXp-concave Optimization (OXO)**, **Online Newton Step (ONS)** achieves minimax optimal regret, but suffers *long runtime* even for simple domains due to *Mahalanobis projections*.
- **LightONS** preserves ONS's regret (even constants) and greatly reduces runtime.
- **For Stochastic eXp-concave Optimization (SXO)**, with online-to-batch conversion, **LightONS** reduces runtime, *answering a COLT'13 open problem*.
- **Importantly**, **LightONS** is "*ONS-like*" and applies to downstream results.
  - Gradient-Norm Adaptivity
  - Joint Efficient Generalized Linear Bandits
  - Memory-Efficient OXO

Thanks!

References

- J. Kivinen, M. Warmuth. *Averaging expert predictions*. **COLT'99**.
- E. Hazan, A. Kalai, S. Kale, A. Agarwal. *Logarithmic regret algorithms for online convex optimization*. **COLT'06**.
- T. Koren. *Open problem: Fast stochastic exp-concave optimization*. **COLT'13**.
- Z. Mhammedi, K. Gatmiry. *Quasi-newton steps for efficient online exp-concave optimization*. **COLT'23**.
- A. Cutkosky, F. Orabona. *Black-box reductions for parameter-free online learning in Banach spaces*. **COLT'18**.
- A. Cutkosky. *Parameter-free, dynamic, and strongly-adaptive online learning*. **ICML'20**.
- Y.J. Zhang, S.A. Xu, P. Zhao, M. Sugiyama. *Generalized linear bandits: Almost optimal regret with one-pass update*. **NeurIPS'25**.
- H. Luo, A. Agarwal, N. Cesa-Bianchi, J. Langford. *Efficient second order online learning by sketching*. **NIPS'16**.
- E. Hazan, T. Koren, K. Levy. *Logistic regression: Tight bounds for stochastic and online optimization*. **COLT'14**.
- D. Foster, S. Kale, H. Luo, M. Mohri, K. Sridharan. *Logistic regression: The importance of being improper*. **COLT'18**.