Chasing Convex Functions with Long-term Constraints: *Optimal Consistency-Robustness Tradeoffs*

Adam Lechowicz
University of Massachusetts Amherst
alechowicz@cs.umass.edu

Nicolas Christianson
California Institute of Technology
nchristianson@caltech.edu

Bo Sun
University of Waterloo
bo.sun@uwaterloo.ca

Noman Bashir
Massachusetts Institute of Technology
nbashir@mit.edu

Mohammad Hajiesmaili
University of Massachusetts Amherst
hajiesmaili@cs.umass.edu

Adam Wierman
California Institute of Technology
adamw@caltech.edu

Prashant Shenoy
University of Massachusetts Amherst
shenoy@cs.umass.edu

ABSTRACT

We introduce and study a family of online metric problems with long-term constraints. In these problems, an online player makes decisions $\mathbf{x}_t$ in a metric space (X, d) to simultaneously minimize their hitting cost $f_t(\mathbf{x}_t)$ and switching cost as determined by the metric. Over the time horizon T, the player must satisfy a long-term demand constraint $\sum_t c(\mathbf{x}_t) \geq 1$, where $c(\mathbf{x}_t)$ denotes the fraction of demand satisfied at time t. Such problems can find a wide array of applications to online resource allocation in sustainable energy/computing systems. We devise optimal competitive and learning-augmented algorithms for specific instantiations of these problems, and further show that our proposed algorithms perform well in numerical experiments.

1 INTRODUCTION

This paper introduces and studies a novel class of online metric problems with *long-term demand constraints*. Our motivation to introduce these problems is rooted in an emerging class of *carbon-aware* control problems for sustainable systems.

This paper builds on a long line of related work that can be roughly classified into two types. In the online metric literature, the problem we study is an extension of *convex function chasing* (CFC) introduced by Friedman and Linial [2], where an online player makes online decisions $\mathbf{x}_t$ in a normed vector space $(X, \|\cdot\|)$ over a sequence of time-varying cost functions in order to minimize their total hitting and switching cost. In the online search literature, the problem we study is a generalization of *one-way trading* (OWT) introduced by El-Yaniv et al. [1], in which an online player must sell an entire asset in fractional shares over a sequence of time-varying prices while maximizing their profit. Existing prior works [3] that simultaneously consider long-term demand constraints (as in OWT) and movement/switching costs (as in CFC) are restricted to unidimensional decision spaces. Generalizing from the unidimensional case is highly non-trivial; e.g., in CFL, the problem cannot simply be decomposed over dimensions due to the shared constraint function and multidimensional switching cost.

In this paper, we obtain results for CFL under problem instantiations that are especially relevant for our motivating applications. We provide tight competitive results for CFL in Theorems 3.3 and 3.4. We propose a learning-augmented algorithm, `CLIP` (Algorithm 1), and show it achieves the provably optimal trade-off between consistency and robustness in Theorems 3.7 and 3.8.

The full version of the paper is available at https://arxiv.org/abs/2402.14012.

2 PROBLEM AND PRELIMINARIES

Convex function chasing with a long-term constraint (CFL). A player chooses decisions $\mathbf{x}_t \in X \subseteq \mathbb{R}^d$ online from a normed vector space $(X, \|\cdot\|)$ in order to minimize their total cost $\sum_{t=1}^{T} f_t(\mathbf{x}_t) + \sum_{t=1}^{T+1} \|\mathbf{x}_t - \mathbf{x}_{t-1}\|$, where $f_t(\cdot) : X \to \mathbb{R}$ is a convex "hitting" cost that is revealed just before the player chooses $\mathbf{x}_t$, and $\|\mathbf{x}_t - \mathbf{x}_{t-1}\|$ is a switching cost associated with changing decisions between rounds. Additionally, the player must satisfy a long term constraint of the form $\sum_{t=1}^{T} c(\mathbf{x}_t) = 1$, where $c(\mathbf{x}) : X \to [0, 1]$ gives the fraction of the constraint satisfied by a decision $\mathbf{x}$. We denote the *utilization* at time t by $z^{(t)} = \sum_{\tau=1}^{t} c(\mathbf{x}_\tau)$ that gives the total fraction of the long-term constraint satisfied up to time t. The offline version of CFL can be formalized as follows:

$$\min_{\{\mathbf{x}_t\}_{t\in[T]}} \underbrace{\sum_{t=1}^{T} f_t(\mathbf{x}_t)}_{\text{Convex hitting cost}} + \underbrace{\sum_{t=1}^{T+1} \|\mathbf{x}_t - \mathbf{x}_{t-1}\|}_{\text{Switching cost}} \text{ s.t. } \underbrace{\sum_{t=1}^{T} c(\mathbf{x}_t) \geq 1,}_{\text{Long-term constraint}}$$
$$\mathbf{x}_t^i \in [0, 1] \ \forall i \in [d], \ \forall t \in [T].$$

Assumptions. We describe the precise variant of CFL for which we design algorithms in the remainder of the paper.

Let $\|\mathbf{x} - \mathbf{x}'\| := \|\mathbf{x} - \mathbf{x}'\|_{\ell_1(\mathbf{w})}$, where $\|\cdot\|_{\ell_1(\mathbf{w})}$ denotes the weighted ℓ_1 norm with weight vector $\mathbf{w} \in \mathbb{R}^d$.

We define the long-term constraint $c(\mathbf{x}) := \|\mathbf{x}\|_{\ell_1(\mathbf{c})}$, i.e., the weighted ℓ_1 norm with weight vector $\mathbf{c} \in \mathbb{R}^d$. Then let the metric space X be the ℓ_1 ball defined by $X := \{\mathbf{x} \in \mathbb{R}^d : c(\mathbf{x}) \leq 1\}$. For all cost functions $f_t(\cdot) : X \to \mathbb{R}$, we assume bounded gradients such that $L \leq {[\nabla f_t]^i}/{\mathbf{c}^i} \leq U \ \forall i \in [d], t \in [T]$, where i denotes the i^{th} dimension of the corresponding vector, and $L, U > 0$ are known constants.

Letting $\mathbf{0}$ denote the origin in $\mathbb{R}^d$ (w.l.o.g), we have the property $f_t(\mathbf{0}) = 0$ for all $t \in [T]$, i.e., that "satisfying none of the long-term constraint costs nothing", since $c(\mathbf{0}) = 0$. We assume the player starts and ends at the origin, i.e., $\mathbf{x}_0 = \mathbf{0}$ and $\mathbf{x}_{T+1} = \mathbf{0}$, to enforce switching "on" and "off." These assumptions are intuitive and reasonable in practice, e.g., in our example motivating application.

Let $\beta := \max\left(\mathbf{w}^i/\mathbf{c}^i\right)$, which gives the greatest magnitude of the switching cost coefficient when normalized by the constraint function. We assume that β is bounded on the interval $[0, {U-L}/{2})$; if β is "too large" (i.e., $> {U-L}/{2}$), the player should prioritize the switching cost.

Recall the player must fully satisfy the long-term constraint before the sequence ends. If the player has satisfied $z^{(t)}$ fraction of it at time t, we assume a compulsory trade begins at time j as soon as $(T - (j+1)) \cdot \mathbf{c}^i < (1 - z^{(j)}) \ \forall i \in [d]$ (i.e., when the time steps after j are not enough to satisfy the constraint). During this compulsory trade, a cost-agnostic algorithm makes maximal decisions to satisfy the constraint. To ensure that the problem remains

Algorithm 1 **C**onsistency **Li**mited **P**seudo-cost minimization (CLIP)

input: consistency parameter ϵ, long-term constraint function $c(\cdot)$, pseudo-cost threshold function $\phi^\epsilon(\cdot)$
initialize: $z^{(0)} = 0$; $p^{(0)} = 0$; $A^{(0)} = 0$; $\texttt{CLIP}_0 = 0$; $\texttt{ADV}_0 = 0$
while cost function $f_t(\cdot)$ is revealed, untrusted advice $\mathbf{a}_t$ is revealed, and $z^{(t-1)} < 1$ **do**
update advice cost $\texttt{ADV}_t = \texttt{ADV}_{t-1} + f_t(\mathbf{a}_t) + \|\mathbf{a}_t - \mathbf{a}_{t-1}\|_{\ell_1(\mathbf{w})}$ and advice utilization $A^{(t)} = A^{(t-1)} + c(\mathbf{a}_t)$
solve *constrained* pseudo-cost minimization problem:

$$\mathbf{x}_t = \underset{\mathbf{x} \in X : c(\mathbf{x}) \le 1 - z^{(t-1)}}{\arg\min} f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - \int_{p^{(t-1)}}^{p^{(t-1)} + c(\mathbf{x})} \phi^\epsilon(u) du$$

$$\text{s.t. } \texttt{CLIP}_{t-1} + f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} + \|\mathbf{x} - \mathbf{a}_t\|_{\ell_1(\mathbf{w})} + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - z^{(t-1)} - c(\mathbf{x}))L + \max((A^{(t)} - z^{(t-1)} - c(\mathbf{x})), 0)(U - L)$$
$$\le (1+\epsilon)[\texttt{ADV}_t + \|\mathbf{a}_t\|_{\ell_1(\mathbf{w})} + (1 - A^{(t)})L]$$

update cost $\texttt{CLIP}_t = \texttt{CLIP}_{t-1} + f_t(\mathbf{x}_t) + \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})}$ and utilization $z^{(t)} = z^{(t-1)} + c(\mathbf{x}_t)$
solve *unconstrained* pseudo-cost minimization problem:

$$\overline{\mathbf{x}}_t = \underset{\mathbf{x} \in X : c(\mathbf{x}) \le 1 - z^{(t-1)}}{\arg\min} f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - \int_{p^{(t-1)}}^{p^{(t-1)} + c(\mathbf{x})} \phi^\epsilon(u) du$$

update pseudo-utilization $p^{(t)} = p^{(t-1)} + \min(c(\overline{\mathbf{x}}_t), c(\mathbf{x}_t))$

technically interesting, we assume that the compulsory trade is a small portion of the sequence.

An example motivating application. Consider a *carbon-aware temporal load shifting* application with heterogeneous servers. Each dimension corresponds to one of d servers. An algorithm makes decisions $\mathbf{x}_t \in \mathbb{R}^d$, where $\mathbf{x}_t^i \in [0,1]$ denotes the load of the ith server at time t. The long-term constraint $\sum_{t=1}^T c(\mathbf{x}_t) \ge 1$ enforces that a job should be finished before time T, and each $\mathbf{c}^i$ represents the throughput of the ith server. Each cost function $f_t(\mathbf{x}_t)$ gives the carbon emissions of servers configured according to $\mathbf{x}_t$, and the switching cost $\|\cdot\|_{\ell_1(\mathbf{w})}$ captures the carbon overhead of reconfiguring the job's allocation.

3 MAIN RESULTS

Competitive algorithm. We build on a generalization of the threshold-based designs used for simple decision spaces in the online search literature called *pseudo-cost minimization*. Our competitive algorithm (ALG1) introduces a novel application of this framework to multidimensional decision spaces, systematically addressing the competitive drawbacks of typical algorithm designs for online metric problems.

Recall that $z^{(t)}$ gives the fraction of the long-term constraint satisfied at time t. We define a function ϕ, which will be used to compute a *pseudo-cost minimization* problem central to ALG1.

Definition 3.1 (Pseudo-cost threshold function ϕ for CFL*). For any utilization $z \in [0,1]$, $\phi(z) = U - \beta + (U/\alpha - U + 2\beta)\exp(z/\alpha)$, where α is the competitive ratio and is defined in (1).*

ALG1 solves the following pseudo-cost minimization problem at each time step. At a high level, the inclusion of ϕ in this pseudo-cost problem enforces that, upon arrival of a cost function, the algorithm satisfies "just enough" of the long-term constraint.

*Definition 3.2 (Pseudo-cost minimization problem (*ALG1*)). At each time step,* ALG1 *solves the following to obtain online decision $\mathbf{x}_t$:*

$$\mathbf{x}_t = \underset{\mathbf{x} \in X : c(\mathbf{x}) \le 1 - z^{(t-1)}}{\arg\min} f_t(\mathbf{x}) + \|\mathbf{x} - \mathbf{x}_{t-1}\|_{\ell_1(\mathbf{w})} - \int_{z^{(t-1)}}^{z^{(t-1)} + c(\mathbf{x})} \phi(u) du$$

In the following, we state that ALG1 is α-competitive and present a lower bound implying α is the best deterministic result for CFL.

Theorem 3.3. ALG1 *is α-competitive for* CFL*, where α is the solution to $\frac{U-L-2\beta}{U-U/\alpha-2\beta} = \exp(1/\alpha)$, given by the following. W is the Lambert W function.*

$$\alpha := \left[W\left(\left(\frac{2\beta}{U} + \frac{L}{U} - 1 \right) e^{\frac{2\beta}{U} - 1} \right) - \frac{2\beta}{U} + 1 \right]^{-1}, \tag{1}$$

Theorem 3.4. *There exists a family of* CFL *instances such that any deterministic online algorithm for* CFL *is at least α-competitive, where α is as defined in (1).*

Learning-augmentation. We present CLIP (**C**onsistency-**Li**mited **P**seudo-cost minimization, Algorithm 1).

Definition 3.5 (Advice model). For CFL *instance $\mathcal{I} \in \Omega$, let* ADV *denote untrusted decision advice, i.e.,* $\texttt{ADV} := \{\mathbf{a}_t \in X : t \in [T]\}$*. If* ADV *is correct, it attains the optimum (i.e.,* $\texttt{ADV}(\mathcal{I}) = \texttt{OPT}(\mathcal{I})$*).*

For $\epsilon \in (0, \alpha - 1]$, we define a *target robustness factor* γ^ϵ, as the unique solution to $\gamma^\epsilon = \epsilon + \frac{U}{L} - \frac{\gamma^\epsilon}{L}(U-L)\ln\left(\frac{U-L-2\beta}{U-U/\gamma^\epsilon - 2\beta}\right)$.

Definition 3.6 (Pseudo-cost threshold function ϕ^ϵ). Given γ^ϵ as above and $p \in [0,1]$, $\phi^\epsilon(p) := U - \beta + (U/\gamma^\epsilon - U + 2\beta)\exp(p/\gamma^\epsilon)$.

For each $t \in [T]$, we define a *pseudo-utilization* $p^{(t)} \in [0,1]$, where $p^{(t)} \le z^{(t)}\ \forall t$, and $p^{(t)}$ describes the "robust satisfaction" of the long-term constraint at time t.

CLIP uses a novel ***projected consistency constraint*** designed to guarantee $(1+\epsilon)$-consistency against ADV by continuously comparing their solutions in terms of the cost incurred so far, the switching cost trajectories, and the projected worst-case cost required to complete the long-term constraint. At a high level, CLIP's *constrained pseudo-cost minimization* yields decisions $\mathbf{x}_t$ that are *as robust as possible* while preserving consistency.

In the following, we state the consistency and robustness results for CLIP and present a lower bound implying that CLIP achieves the optimal tradeoff between the two.

Theorem 3.7. *For any $\epsilon \in [0, \alpha - 1]$,* CLIP *is $(1+\epsilon)$-consistent and γ^ϵ-robust for* CFL *(γ^ϵ as defined above).*

Theorem 3.8. *Given untrusted advice* ADV *and $\epsilon \in (0, \alpha - 1]$, any $(1+\epsilon)$-consistent learning-augmented algorithm for* CFL *is at least γ^ϵ-robust, where γ^ϵ is defined above.*

REFERENCES

[1] Ran El-Yaniv, Amos Fiat, Richard M. Karp, and G. Turpin. 2001. Optimal Search and One-Way Trading Online Algorithms. *Algorithmica* 30, 1 (May 2001), 101–139.
[2] Joel Friedman and Nathan Linial. 1993. On convex body chasing. *Discrete & Computational Geometry* 9, 3 (March 1993), 293–321.
[3] Adam Lechowicz, Nicolas Christianson, Bo Sun, Noman Bashir, Mohammad Hajiesmaili, Adam Wierman, and Prashant Shenoy. 2024. Online Conversion with Switching Costs: Robust and Learning-augmented Algorithms. arXiv:2310.20598 [cs.DS] https://arxiv.org/abs/2310.20598