

ADVERSARIAL MULTI-AGENT TARGET TRACKING WITH INEXACT ONLINE GRADIENT DESCENT

Amrit Singh Bedi, Paban Sarma, and Ketan Rajawat

Department of Electrical Engineering, IIT Kanpur, U.P., India
{amritbd@iitk.ac.in, sarma.paban@gmail.com, ketan@iitk.ac.in}

ABSTRACT

Multi-agent systems are being increasingly deployed in challenging environments for performing complex tasks such as multi-target tracking, search-and-rescue, and intrusion detection. This paper formulates the generic target tracking problem as a time-varying optimization problem and puts forth an inexact online gradient descent method for solving it sequentially. The performance of the proposed algorithm is studied by characterizing its dynamic regret, a notion common to the online learning literature. Building upon the existing results, we provide improved regret rates that not only allow non-strongly convex costs but also explicating the role of the cumulative gradient error. The objective function is convex but the variable belongs to a compact domain. The efficacy of the proposed inexact gradient framework is established on a multi-agent multi-target tracking problem.

Index terms: Time varying optimization, stochastic optimization, target tracking, gradient descent methods.

1. INTRODUCTION

Multi-agent systems are increasingly being used for challenging tasks such as multi-target tracking [1, 2], planetary exploration and mapping [3], search-and-rescue, and intrusion detection [4]. Achieving such team-based goals requires the robotic platforms to not only sense and understand the environment, but also cooperate among themselves through judicious data exchange and fusion. Consequently, resource allocation and optimization becomes an important aspect of the overall motion planning problem. Indeed, the recent trend is to formulate target tracking as a constrained convex optimization problem that must be solved at every time step [5–9]. Such time-varying optimization problems have their origins in the control theory literature, where they have been applied to path planning [10] and dynamic parameter tracking problems [11, 12]. Given the limited computational and communications capabilities of the mobile robots, solving the full per-time instant optimization problem before taking the action may not necessarily be viable. Instead, recent works have advocated the use of simpler one-iteration algorithms such as the online interior point, prediction-correction, online ADMM, and gradient descent methods, that have been shown to approach the optimal asymptotically. Leveraging the tools from classical optimization theory, these dynamic optimization algorithms not only admit low-complexity distributed implementations, but are also amenable to analytical performance guarantees.

Online machine learning represents a parallel but closely related development that has been widely applied to solve problems in Big Data [13]. First proposed by [5], the online convex optimization framework models the agents as learners and targets as adversaries. Within this sequential learning paradigm, the learner performs an action and the adversary reveals a corresponding loss function at each

time step. The eventual goal of the learner is to minimize the cumulative loss. Recent year have witnessed the development of theoretical guarantees in form of *dynamic regret*, where the performance of the learner is measured against that of an adaptive and time-varying adversary [8, 14, 15].

1.1. Related work and contributions

Inexact gradient methods have been widely used to solve a variety of optimization problems, especially in the context of machine learning [16, 17]. Since calculating an approximate gradient is often cheap, recent years have witnessed the development of several variants, such as the incremental aggregated gradient method [18], stochastic average gradient method [19], stochastic variance reduced gradient method [20], SAGA [21]. For static optimization problems, the inexact gradient methods are known to converge at a linear rate even for non-strongly convex objectives [22], [23]. For time varying optimization problem, dynamic regret was first introduced in [5]. As compared to the weaker notion of static regret, the idea here is to compare the performance of the tracker against that of an adaptive and time-varying adversary [8, 9, 14]. Dynamic regret bounds for the gradient descent and related first order methods have been reported in [5, 6, 8, 9, 14, 24, 25]. As compared to existing results, the bounds provided here are not only stronger but also require minimal assumptions on the cost function.

This paper studies the multi-agent multi-target tracking problem from the lens of online convex optimization. Prompted by the noisy and possibly incorrect target position information available to the agents, we put forth the inexact online gradient descent (IOGD) algorithm. The key theoretical contribution is the development of improved dynamic regret bounds for non-strongly convex objective functions. Improving the existing results for general convex problems, it is shown that the dynamic regret is bounded by the *path length* of the target [See [26] for proofs]. Different from the existing literature, we also explicate the dependence of the dynamic regret bounds on the cumulative gradient errors, that are otherwise allowed to be adversarial. Finally, the flexibility of the IOGD algorithm is demonstrated by applying it to the multi-agent multi-target tracking problem from [27].

2. PROBLEM FORMULATION

We consider the general problem of tracking a time-varying parameter that evolves according to an unknown dynamic model. The general setting considered here subsumes the target tracking application of interest, where the parameter may represent the location(s) of the target(s) being pursued by one or more agents. As motivated in [8, 15, 28], the parameter at time k can be written as the solution of

Table 1: Summary of related works on time-varying optimization (cf. Sec.3)

References	Loss function	Inexact	Function	Regret rate
[5]	Convex	No	Deterministic	$\mathcal{O}(\sqrt{K}(1 + W_K))$
[14]	Convex	No	Stochastic	$\mathcal{O}(K^{2/3}(1 + V_K)^{1/3})$
[24]	Convex	No	Deterministic	$\mathcal{O}(\sqrt{D_K + 1} + \min\{\sqrt{(D_K + 1)W_K}, [(D_K + 1)V_K K]^{1/3}\})$
[25]	Convex	No	Stochastic	$\mathcal{O}(\sqrt{KW'_K})$
[9]	Convex+QG	No	Deterministic	$\mathcal{O}(1 + W'_K)$
This work	Convex (P ₂)	Yes	Deterministic/Stochastic (A1)	$\mathcal{O}(1 + \sqrt{KE_K} + W_K)$
This work	Convex (P ₂)	Yes	Stochastic (A2)	$\mathcal{O}(1 + E_K + W_K)$

the following (discrete) time-varying convex optimization problem

$$\mathbf{x}_k^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} f_k(\mathbf{x}) \quad k = 1, 2, \dots \quad (1)$$

where f_k is a smooth convex function and $\mathcal{X} \subset \mathbb{R}^n$ is a compact convex set with diameter R . The set notation in (1) emphasizes the fact that in general, the minimizer of f_k may not necessarily be unique. The parameter estimate at time k is denoted by \mathbf{x}_k . The agents do not know the full functional form of f_k but are only revealed an inexact version of the gradient $\tilde{\nabla} f(\mathbf{x}_k) := \nabla f_k(\mathbf{x}_k) + \mathbf{e}_k$ for some $\mathbf{e}_k \in \mathbb{R}^n$. The agents make use of these inexact gradients to improve their estimates of \mathbf{x}_k^* in an online manner.

This paper considers the problem from an online convex optimization perspective, viewing the agents as learners and targets as adversaries. Specifically, at time k , the online learner selects an action $\mathbf{x}_k \in \mathcal{X}$ and incurs a cost $f_k(\mathbf{x}_k)$, where $f_k : \mathbb{R}^n \rightarrow \mathbb{R}$ are smooth convex functions selected by the adversary. In response to the agent's action, the adversary also reveals an inexact gradient $\tilde{\nabla} f_k(\mathbf{x}_k)$ to the learner. The goal of the learner is to minimize its cumulative loss $\sum_{k=1}^K f_k(\mathbf{x}_k)$ over K time slots. Of particular interest is to characterize the so-called *dynamic regret* of the learner, that measures the cumulative mismatch between the learner's action and the optimal action [6, 8, 14]:

$$\mathbf{Reg}_K := \sum_{k=1}^K f_k(\mathbf{x}_k) - f_k(\mathbf{x}_k^*) \quad (2)$$

where \mathbf{x}_k^* is as defined in (1). In order for the tracking to be successful, the dynamic regret must be upper bounded by a sublinear function of K .

2.1. Parameter variations and error bounds

It is well known that a sublinear dynamic regret may not always be achievable, e.g., if the parameter variations or the gradient errors are too large [9, 14]. The goal of the current paper will therefore be to bound the dynamic regret using functions of the cumulative parameter variations and errors. For the target tracking setting at hand, it makes sense to characterize the parameter variations using the path length, defined as

$$W_K := \sum_{k=2}^K \|\mathbf{x}_k^* - \mathbf{x}_{k-1}^*\| \quad (3)$$

for some sequence of parameter values $\{\mathbf{x}_k^*\}_{k \geq 1}$. More generally there exist a class of related complexity measures that can be used to characterize the parameter variations [8].

The gradient errors \mathbf{e}_k can be modeled either as being deterministic with bounded norms or stochastic with bounded variance. Deterministic errors may be of interest in adversarial settings while stochastic errors are useful for modeling communication and computational noise. In order to unify the subsequent development, a

generic stochastic error model is considered that subsumes that deterministic case. Let \mathcal{F}_k denote the sigma field generated by the random sequence $\{\mathbf{e}_\tau\}_{\tau=1}^{k-1}$. The following assumption bounds the second moment of the error.

A1. Error bound: *The stochastic sequence $\{\mathbf{e}_k\}_{k=1}^K$ adheres to the following bound on the second moment:*

$$\mathbb{E} [\|\mathbf{e}_k\|^2 \mid \mathcal{F}_k] \leq \varepsilon_k^2 + \nu^2 \|\nabla f_k(\mathbf{x}_k)\|^2 \quad (4)$$

with $\varepsilon_k \leq \varepsilon < \infty$ for all $k \geq 1$, and $\nu \geq 0$ and $\varepsilon \geq 0$ are constants.

When the errors are deterministic, (4) is equivalent to a worst-case bound on the error norm. In the general case, the goal will be to establish the sublinearity of the expected dynamic regret $\mathbb{E} [\mathbf{Reg}_K]$. The specific form of the bounds in (A1) is inspired from [29, 30] and allows for errors that are proportional to the gradient norm in addition to an additive term. It is remarked that from Jensen's inequality, (4) implies that $\mathbb{E} [\|\mathbf{e}_k\| \mid \mathcal{F}_k] \leq \varepsilon_k + \nu \|\nabla f_k(\mathbf{x}_k)\|$. The required dynamic regret bounds will be developed in terms of the path length W_K and the cumulative error bound $E_K := \sum_{k=1}^K \varepsilon_k$.

A particular case of interest is when the gradient errors constitute a white noise process as specified in the following assumption.

A2. White noise: *The zero-mean error sequence $\{\mathbf{e}_k\}_{k=1}^K$ is independent identically distributed, i.e.,*

$$\mathbb{E} [\mathbf{e}_k \mid \mathcal{F}_k] = \mathbf{0}. \quad (5)$$

Assumption (A2) may be applicable, for instance, when the gradient errors arise from communication errors. The requirement in (5) is more restrictive than that in (4), but also results in improved regret bounds.

Remark 1. The path length definition used in (3) applies to an arbitrary sequence of true parameter values $\{\mathbf{x}_k^*\}$ and does not depend on \mathcal{X}_k^* . Consequently, the path length definition in (3) is stronger than those used in [9, 25]. In particular, the definitions in [9, 25] take the following form

$$W'_K := \max_{\{u_k \in \mathcal{X}_k^*\}_{k=2}^K} \sum_{k=2}^K \|\mathbf{u}_k - \mathbf{u}_{k-1}\| \quad (6)$$

$$W''_K := \sum_{k=2}^K \max_{\mathbf{u} \in \mathcal{X}_k^*} \|P_k(\mathbf{u}) - P_{k-1}(\mathbf{u})\|, \quad (7)$$

respectively, where $P_k(\mathbf{u}) := \arg \min_{\mathbf{y} \in \mathcal{X}_k^*} \|\mathbf{y} - \mathbf{u}\|$. The use of an arbitrary trajectory in (3) allows us handle such unbounded sets while also ensuring that $W_K \leq W'_K$ and $W_K \leq W''_K$ for problems where \mathcal{X}_k^* is compact.

3. PROPOSED ALGORITHM AND ASSUMPTIONS

The online gradient descent algorithm has been widely used to solve online learning problems owing to its flexibility and simplicity [5, 6,

8, 14, 24]. This work considers the inexact online gradient descent (IOGD) method that takes the form:

$$\mathbf{x}_{k+1} = P_{\mathcal{X}}[\mathbf{x}_k - \alpha(\nabla f_k(\mathbf{x}_k) + \mathbf{e}_k)] \quad (8)$$

where $P_{\mathcal{X}}(\cdot)$ denotes the projection onto the set \mathcal{X} . The IOGD method has also been applied to static and online problems [22]. The IOGD method is also closely related to the incremental and variance reduced variants of the gradient descent algorithm. The full algorithm is summarized in Algorithm 1.

Algorithm 1 IOGD: Inexact Online Gradient Descent

- 1: Initialize \mathbf{x}_1
 - 2: **for** $k = 1, 2, \dots$ **do**
 - 3: **Action** \mathbf{x}_k
 - 4: **Observe** inexact gradient $\nabla f_k(\mathbf{x}_k) + \mathbf{e}_k$
 - 5: **Update** $\mathbf{x}_{k+1} = P_{\mathcal{X}}[\mathbf{x}_k - \alpha(\nabla f_k(\mathbf{x}_k) + \mathbf{e}_k)]$
 - 6: **end for**
-

In addition to Assumptions (A1) and (A2) stated in Sec. 2, the subsequent analysis will also require the following regularity conditions.

A3. Lipschitz continuity: The function ∇f_k is Lipschitz continuous with parameter L :

$$\|\nabla f_k(\mathbf{u}) - \nabla f_k(\mathbf{v})\| \leq L \|\mathbf{u} - \mathbf{v}\| \quad (9)$$

for all $k \geq 1$ and $\mathbf{u}, \mathbf{v} \in \mathcal{X}$.

A4. Vanishing gradient: The optimum \mathbf{x}_k^* lies in the relative interior of the set \mathcal{X} , that is, $\nabla f_k(\mathbf{x}_k^*) = 0$ for all $\mathbf{x}_k^* \in \mathcal{X}_k^*$.

A5. Bounded Variation: For a given \mathbf{x}_k^* , there exists some $\sigma > 0$ such that $\|\mathbf{x}_{k+1}^* - \mathbf{x}_k^*\| \leq \sigma$ for all $k \in \mathbb{N}$.

Of these, both (A3) and (A4) are standard and apply to large class of online learning problems. Likewise, the requirement in (A5) imposes a limit on the maximum velocity of the target and is therefore applicable to most target tracking problems. The bounded variation condition is also satisfied, for instance, if W_K is sublinear or linear and the target motion is not too ‘jumpy’.

4. REGRET BOUNDS

The results would be presented here for two specific scenarios, namely (a) gradient errors following (A1) with $\nu = 0$, and (b) gradient errors following (A1)-(A2) but possibly non-zero value of ν . We begin with stating the following intermediate lemma

Lemma 1. Under (A1) with $\nu = 0$, (A3)-(A4), and for a sequence $\{\mathbf{x}_k^*\}$ satisfying (3), the IOGD iterates satisfy

$$\mathbb{E}[\|\mathbf{x}_{k+1} - \mathbf{x}_k^*\|] \leq \mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_k^*\|] + \frac{\xi}{R}(\mathbb{E}[f_k(\mathbf{x}_k) - f_k(\mathbf{x}_k^*)]) + s_k$$

where $\xi := 1 - 2\alpha(1 - 2\alpha L)$ and $s_k := \sqrt{2\alpha^2 \varepsilon_k^2 + 2\alpha \varepsilon_k R}$.

Lemma 1 leads directly to the required dynamic regret bounds under (A1) with $\nu = 0$. This lemma establishes that the distance of next action \mathbf{x}_{k+1} from the current optimal \mathbf{x}_k^* is upper bounded by the quantities calculated at current time k .

Theorem 1. Under (A1) with $\nu = 0$, (A3)-(A4), and for a sequence $\{\mathbf{x}_k^*\}$ satisfying (3), the IOGD iterates adhere to the following dynamic regret rate

$$\mathbb{E}[\mathbf{Reg}_K] \leq \mathcal{O}(1 + \sqrt{K E_K} + W_K). \quad (10)$$

This theorem establishes that the proposed algorithm provides sublinear regret when E_K and W_K are sublinear. This results states that the proposed algorithm will become exactly close to optimal solution for large enough K . The results can be improved for the case when the gradient errors follow a white noise process as provided in following corollary.

Corollary 1. Under (A1)-(A4), and for a sequence $\{\mathbf{x}_k^*\}$ satisfying (3), the IOGD iterates satisfy

$$\mathbb{E}[\|\mathbf{x}_{k+1} - \mathbf{x}_k^*\|] \leq \mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_k^*\|] + \frac{\xi_2}{R}(\mathbb{E}[f_k(\mathbf{x}_k) - f_k(\mathbf{x}_k^*)]) + \varepsilon_k \quad (11)$$

where $\xi_2 := 1 - 2\alpha(1 - (1 + \nu^2)\alpha L)$. For this case, the dynamic regret is bounded as

$$\mathbb{E}[\mathbf{Reg}_K] \leq \mathcal{O}(1 + E_K + W_K). \quad (12)$$

It is remarked that there are various existing online algorithms which can be obtained as special case of the proposed inexact online gradient descent algorithm in this paper. For example, Incremental OGD with increasing sample size, and Proximal OGD methods etc (See [26] for details).

5. NUMERICAL TEST: MULTI-TARGET TRACKING

This section develops a low-complexity online multi-target tracking algorithm inspired from the convex optimization-based target tracking framework developed in [1]. Specifically, a team of n agents at locations $\{\mathbf{x}_k^i\}_{i=1}^n$ is tasked with tracking a set of m targets at locations $\{\mathbf{y}_k^j\}_{j=1}^m$. The discretized problem is formulated as the following convex optimization that must be solved for each $k \geq 1$ [1, Thm. 3.8.1]:

$$\{\mathbf{x}_{k+1}^i\} = \arg \min_{\{\mathbf{x}^i \in \mathbb{R}^p\}} \sum_{i=1}^n \psi_k^i(\mathbf{x}^i) \quad (13a)$$

$$\text{s.t. } \|\mathbf{x}^i - \mathbf{x}_k^i\|^2 \leq v_i = 1, \dots, n \quad (13b)$$

$$\sum_{i=1}^n w_k^{ij} \|\mathbf{x}^i - \mathbf{y}_k^j\|^2 \leq \eta_j = 1, \dots, m \quad (13c)$$

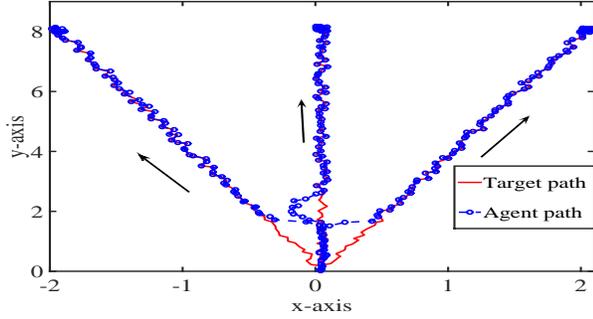
where $\psi_k^i(\cdot)$ is a time-varying cost function and v is the square of the maximum distance that an agent may cover within a single time slot. For this paper, the following cost function is used:

$$\psi_k^i(\{\mathbf{x}^i\}) = \frac{1}{2} \sum_{\ell \neq i} \|\mathbf{x}^i - \mathbf{x}^\ell\|^2 + \gamma \|\mathbf{x}^i - \mathbf{x}_k^i\|^2 \quad (14)$$

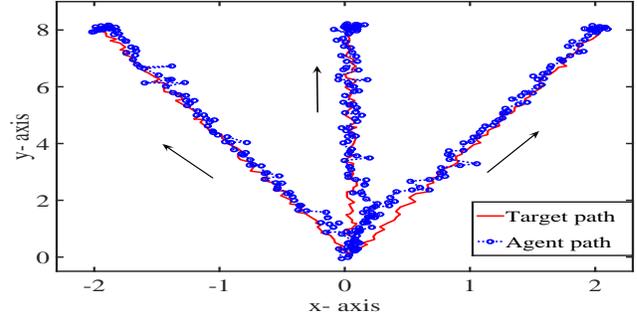
where $\gamma > 0$ is a regularization parameter. The objective function encourages agent i to remain close to the other agents. At the same time, the regularization term forces the agents to not move around unnecessarily. The constraint in (13c) is the linearized version of the original constraint obtained from using the process described in [1, Chap. 3]. A sigmoidal weight function is utilized that takes the form:

$$w_k^{ij} = \left(1 + e^{-\omega(\epsilon - \|\mathbf{x}_k^i - \mathbf{y}_k^j\|)}\right)^{-1} \quad (15)$$

where ω and ϵ are positive parameters. Observe that the weights are small for agent-target pairs that are far from each other. In other words, the constraint in (13c) encourages the set of agents tracking a target j to stay close to it. The weights are normalized so that $\sum_{i=1}^n w_k^{ij} = 1$ in order to ensure that each target is tracked by at



(a) Target tacking with $m = 3$ and $n = 3$ (Proposed algorithm).



(b) Target tacking with $m = 3$ and $n = 3$ (algorithm in [27]).

Fig. 1: Real time target tracking with proposed algorithm and algorithm of [27].

least one agent. Finally, it is remarked that the agents may only know the estimated target locations $\{\hat{\mathbf{y}}_k^j\}_{j=1}^m$ instead of the true locations required for (13c). Since the constrained optimization problem in (13) is not of the form required in (1), the IOGD algorithm will instead be applied in the dual domain. It can be verified that (13) satisfies the Slater's conditions, and therefore has zero duality gap. To this end, associate dual variables $\{\lambda^i\}_{i=1}^n$ and $\{\nu^j\}_{j=1}^m$ with (13b) and (13c), respectively. Collecting the primal and dual variables $\{\lambda^i\}$, $\{\nu^j\}$, and $\{\mathbf{x}^i\}$ into vectors $\boldsymbol{\lambda}$, $\boldsymbol{\nu}$, and \mathbf{x} respectively, the Lagrangian can be written as

$$L_k(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = \sum_{i=1}^n \sum_{\ell=i+1}^n \|\mathbf{x}^i - \mathbf{x}^\ell\|^2 + \gamma \|\mathbf{x} - \mathbf{x}_k\|^2 \quad (16)$$

$$+ \sum_{i=1}^n \lambda^i (\|\mathbf{x}^i - \mathbf{x}_k^i\|^2 - v) + \sum_{j=1}^m \nu^j \left[\sum_{i=1}^n w_k^{ij} \|\mathbf{x}^i - \mathbf{y}_k^j\|^2 - \eta \right].$$

Thus the dual function can be written as

$$\varrho_k(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \arg \min_{\mathbf{x}} L_k(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) \quad (17)$$

Since the dual function is always concave, the proposed IOGD algorithm can be utilized to maximize ϱ_k in an online fashion. Then solving the dual problem using the proposed technique results in algorithm 2.

Algorithm 2 IOGD-based multi-target tracking

- 1: **Initialize** \mathbf{x}_1 , $\boldsymbol{\lambda}_1$, and $\boldsymbol{\nu}_1$, and step sizes α_λ , and α_ν
- Repeat for** $k = 1, 2, \dots$,
- 2: Compute weights $\{w_k^{ij}\}$ from (15)
- 3: **Calculate** the next location as
 $\mathbf{x}_{k+1} = \arg \min_{\mathbf{x}} L_k(\mathbf{x}, \boldsymbol{\lambda}_k, \boldsymbol{\nu}_k)$
- 4: Update for all agents and targets:

$$\lambda_{k+1}^i = \left[\lambda_k^i - \alpha_\lambda (\|\mathbf{x}_{k+1}^i - \mathbf{x}_k^i\|^2 - v) \right]_+ \quad (18)$$

$$\nu_{k+1}^j = \left[\nu_k^j - \alpha_\nu \left(\sum_{i=1}^n w_k^{ij} \|\mathbf{x}_{k+1}^i - \hat{\mathbf{y}}_k^j\|^2 - \eta \right) \right]_+ \quad (19)$$

The performance of the proposed multi-target tracking algorithm is studied on a number of simulated planar environments. The agent velocities are restricted to 10m/s and a target is assumed covered if it is within one meter from the agent. As a toy example,

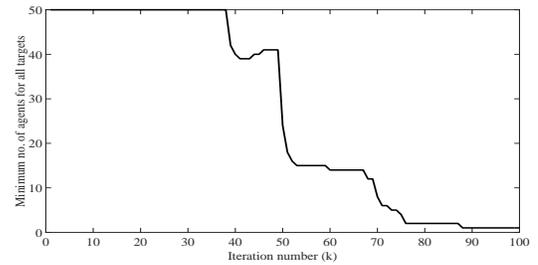


Fig. 2: Number of agents covering each target ($m = 10$, $n = 50$).

consider first a simple scenario consisting of three targets ($m = 3$) and three agents $n = 3$. The targets are co-located at time $k = 1$ and start moving away from each other along the paths shown in Fig. 1a. It can be seen that the proposed algorithm works as expected, and the agent team splits up in order to track the three targets. On the other hand, the algorithm in [27] does not necessarily exhibit such a behavior and requires careful parameter tuning so as to allow tracking with reasonable accuracy; see Fig. 1b. Indeed, since [27] entailed solving a constrained convex optimization problem at every time instant, it was observed that unless the parameters are not selected carefully, the problem could become infeasible. It was however possible to circumvent this behavior to a certain extent by explicitly adding noise to the output of the optimization problem. In contrast, no such issue was present in the proposed IOGD algorithm, whose performance was quite robust to the choice of parameters.

Next, we consider a large scale system with $m = 10$ targets and $n = 50$ agents. As expected, the IOGD algorithm is capable of tracking most of the targets at low complexity. As with the smaller system considered earlier, the splitting of the agent teams is observed and is evident from the supplementary video included with this paper¹. It is important to emphasize that the tracking performance of the IOGD is at par with the convex optimization approach of [27]. In contrast, solving a general convex optimization problem as required in [27] incurs a complexity of at least $\mathcal{O}(n^3)$ as opposed to the $\mathcal{O}(n)$ complexity incurred in the calculation of the inexact gradient in (18)-(19). For the sake of comparison, both algorithms were implemented in MATLAB and their run-times measured on an Intel Xeon E3-1226 3.30GHz CPU machine. The resulting per-iteration run-time for the proposed algorithm was 49 ms, as compared to that of 974 ms required by [27].

¹<https://www.youtube.com/watch?v=bVto6LIteHM>

6. REFERENCES

- [1] Jason C Derenick, *A Convex Optimization Framework for Multi-agent Motion Planning*, ProQuest, 2009.
- [2] Sheng-Yuan Tu and Ali H Sayed, "On the influence of informed agents on learning and adaptation over networks," *IEEE Trans. Signal Process.*, vol. 61, no. 6, pp. 1339–1356, 2013.
- [3] A. Trafton, "Team envisions exploring mars with mini probes (MIT news office)."
- [4] Pieter Abbeel, Adam Coates, Morgan Quigley, and Andrew Y Ng, "An application of reinforcement learning to aerobic helicopter flight," in *Adv. Neural Inf. Process. Syst.*, 2007, pp. 1–8.
- [5] Martin Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proc. of the 20th Int. Conf. on Machine Learning*, 2003, pp. 928–936.
- [6] Eric C Hall and Rebecca M Willett, "Online convex optimization in dynamic environments," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 4, pp. 647–662, 2015.
- [7] Andrea Simonetto, Alec Koppel, Aryan Mokhtari, Geert Leus, and Alejandro Ribeiro, "Decentralized prediction-correction methods for networked time-varying convex optimization," *IEEE Trans. on Autom. Control*, 2017.
- [8] Aryan Mokhtari, Shahin Shahrampour, Ali Jadbabaie, and Alejandro Ribeiro, "Online optimization in dynamic environments: Improved regret rates for strongly convex problems," in *55th IEEE Conf. on Decision and Control (CDC)*, 2016, pp. 7195–7201.
- [9] Lijun Zhang, Tianbao Yang, Jinfeng Yi, Jing Rong, and Zhi-Hua Zhou, "Improved dynamic regret for non-degenerate functions," in *Adv. Neural Inf. Process. Syst.*, 2017, pp. 732–741.
- [10] CL Castillo, W Moreno, and KP Valavanis, "Unmanned helicopter waypoint trajectory tracking using model predictive control," in *IEEE Mediterranean Conference on Control & Automation (MED'07)*, 2007, pp. 1–8.
- [11] Reza Abdoolee, Benoit Champagne, and Ali H Sayed, "Estimation of space-time varying parameters using a diffusion lms algorithm," *IEEE Trans. Signal Process.*, vol. 62, no. 2, pp. 403–418, 2014.
- [12] Ke Zhou and Stergios I Roumeliotis, "Multirobot active target tracking with combinations of relative observations," *IEEE Trans. Robot.*, vol. 27, no. 4, pp. 678–695, 2011.
- [13] Dilhan Perera, Judy Kay, Irena Koprinska, Kalina Yacef, and Osmar R Zaiane, "Clustering and sequential pattern mining of online collaborative learning data," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 6, pp. 759–772, 2009.
- [14] Omar Besbes, Yonatan Gur, and Assaf Zeevi, "Non-stationary stochastic optimization," *Operations Research*, vol. 63, no. 5, pp. 1227–1244, 2015.
- [15] S. Shahrampour and A. Jadbabaie, "Distributed online optimization in dynamic environments using mirror descent," *IEEE Trans. Autom. Control*, vol. PP, no. 99, pp. 1–1, 2017.
- [16] Hui Zhang and Lizhi Cheng, "Restricted strong convexity and its applications to convergence analysis of gradient-type methods in convex optimization," *Optimization Letters*, vol. 9, no. 5, pp. 961–979, 2015.
- [17] Michael P Friedlander and Mark Schmidt, "Hybrid deterministic-stochastic methods for data fitting," *SIAM J. Sci. Comput.*, vol. 34, no. 3, pp. A1380–A1405, 2012.
- [18] Doron Blatt, Alfred O Hero, and Hillel Gauchman, "A convergent incremental gradient method with a constant step size," *SIAM J. on Opt.*, vol. 18, no. 1, pp. 29–51, 2007.
- [19] Nicolas L Roux, Mark Schmidt, and Francis R Bach, "A stochastic gradient method with an exponential convergence rate for finite training sets," in *Adv. Neural Inf. Process. Syst.*, 2012, pp. 2663–2671.
- [20] Rie Johnson and Tong Zhang, "Accelerating stochastic gradient descent using predictive variance reduction," in *Adv. Neural Inf. Process. Syst.*, 2013, pp. 315–323.
- [21] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien, "SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives," in *Adv. Neural Inf. Process. Syst.*, 2014, pp. 1646–1654.
- [22] Anthony Man-Cho So and Zirui Zhou, "Non-asymptotic convergence analysis of inexact gradient methods for machine learning without strong convexity," *Optimization Methods and Software*, vol. 32, no. 4, pp. 963–992, 2017.
- [23] Zhi-Quan Luo and Paul Tseng, "On the linear convergence of descent methods for convex essentially smooth minimization," *SIAM J. Control and Opt.*, vol. 30, no. 2, pp. 408–425, 1992.
- [24] Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan, "Online optimization: Competing with dynamic comparators," in *Artificial Intelligence and Statistics*, 2015, pp. 398–406.
- [25] Tianbao Yang, Lijun Zhang, Rong Jin, and Jinfeng Yi, "Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient," in *International Conference on Machine Learning*, 2016, pp. 449–457.
- [26] Amrit Singh Bedi, Paban Sarma, and Ketan Rajawat, "Adversarial multi-agent target tracking with inexact online gradient descent," *arXiv preprint arXiv:1710.05133*, 2017.
- [27] Jason Derenick, John Spletzer, and Ani Hsieh, "An optimal approach to collaborative target tracking with performance guarantees," *J. Intell. and Robotic Syst.*, vol. 56, no. 1-2, pp. 47–67, 2009.
- [28] Andrea Simonetto and Geert Leus, "On non-differentiable time-varying optimization," in *IEEE 6th Int. Workshop on Computational Advances in Multi-Sensor Adaptive Process.*, 2015, pp. 505–508.
- [29] Dimitri P Bertsekas and John N Tsitsiklis, "Gradient convergence in gradient methods with errors," *SIAM J. on Opt.*, vol. 10, no. 3, pp. 627–642, 2000.
- [30] Lorenzo Rosasco, Silvia Villa, and Bang Công Vũ, "Convergence of stochastic proximal gradient algorithm," *arXiv preprint arXiv:1403.5074*, 2014.