Online Utility-Optimal Trajectory Design for Time-Varying Ocean Environments

Mohan Krishna Nutalapati, *Student Member, IEEE*, Shruti Joshi and Ketan Rajawat, *Member, IEEE*

***Abstract*— This paper considers the problem of online optimal trajectory design under time-varying environments. Of particular interest is the design of energy-efficient trajectories under strong and uncertain disturbances in ocean environments and time-varying goal location. We formulate the problem within the constrained online convex optimization formalism, and a modified online gradient descent algorithm is motivated. The mobility constraints are met using a carefully chosen step-size, and the proposed algorithm is shown to incur sublinear regret. Different from the state-of-the-art algorithms that entail planning and re-planning the full trajectory using forecast data at each time instant, the proposed algorithm is entirely online and relies mostly on the current ocean velocity measurements at the vehicle locations. The trade-off between excess delay incurred in reaching the goal and the overall energy consumption is examined via numerical tests carried out on real data obtained from the regional ocean modelling system. As compared to the state-of-the-art algorithms, the proposed algorithm is not only energy-efficient but also several orders of magnitude computationally efficient.**

I. Introduction

Trajectory design for motion planning is one of the core building blocks of autonomous systems. The foremost goal here is to design optimal trajectories starting and ending at specified locations while satisfying application-oriented constraints such as (i) avoiding clutter and other obstacles [1]; (ii) maintaining connectivity [2]; (iii) adhering vehicle-specific constraints such as turn radii or spatial envelop of the vehicle [3]; and (iv) energy efficiency [4]. Energy efficiency is critical in naval and aerial environments where the vehicle may simply drift along the surrounding air/water flow. Designing trajectories that are energy efficient is important, for instance, in oceanic environments vehicles with constrained energy budget are often deployed for long-term autonomous missions [5] such as surveying, mine hunting, chasing seaborne targets, oceanographic research etc.

Minimum energy trajectory design can generally be formulated as a convex optimization problem with intermediate waypoint locations as variables. The objective function encourages lower energy consumption while the constraints may impose physical restrictions arising from maximum sustainable velocity, smoothness requirements, etc., [4]. In practice, a maximum time of arrival is also specified to rule out highly energy efficient trajectories that may take the vehicle too far from the straight line path. Indeed, a trade-off exists between the energy consumption and the excess delay incurred in reaching the goal [6], [7]. This work considers the problem of designing energy optimal trajectories for a given excess delay.

The authors are with the Department of Electrical Engineering, Indian Institute of Technology Kanpur, Kanpur 208016, India, (e-mail: nmohank@iitk.ac.in; shrutij@iitk.ac.in; ketan@iitk.ac.in).

Classical optimal trajectory design approaches require information of ocean currents between the start and the goal locations [8]. However, unless the trajectories are short, such information is only available in the form of low-resolution historical data or forecasts. For instance, approximate forecasts available from Regional Ocean Modeling System (ROMS) [9] or Navy Coastal Ocean Model [10] provide ocean current velocities at points separated in the order of kilometers. Further, since the forecasts are obtained from satellite or field observations, they are available either at very coarse resolutions or only for specific geographical regions. Owing to the uncertainty in the ocean currents forecasts, offline trajectory design of autonomous vehicles is both challenging and potentially dangerous due to the possibility of being swept into shipping lanes or land [11]. To counter these challenges, state-of-the-art approaches rely on a re-planning framework using the forecast disturbances with predictions from Gaussian process (GP) regression added to it [4]. It is remarked that although such algorithms are capable of handling time-varying environments and a moving goal, the reliance on forecasts and heavy computational costs make them impractical for many systems.

This work puts forth an online algorithm for utility-optimal trajectory design in unknown and time-varying environments. Following the spirit of online convex optimization, the proposed algorithm generates the subsequent waypoint locations in an online fashion and without relying on forecasts. The proposed online gradient descent (OGD) algorithm builds upon the utility-optimal trajectory design framework in the context of communication networks and incorporates information about the excess delay budget within the design. Since each update step only uses the information at the current time instant, both the ocean current velocities as well as the goal locations are allowed to be time-varying. Unlike the sampling-based algorithms used in [12] and [4] whose outputs are random, the proposed algorithm does not include any random parameters and always produces the same trajectory for a given system. Finally, the trade-off between the total energy consumption and excess delay is empirically evaluated using ocean currents data collected from [9].

In summary, the main contributions of the current work include (a) development and analysis of the OGD algorithm for generic time-varying utility optimal trajectory design problems, yielding a sublinear regret; (b) formulation of the online trajectory design problem for a watercraft or an unmanned surface vehicles (USV) operating under strong disturbances; and (c) demonstration of the low complexity and superior performance of the proposed algorithm as compared to the significantly more sophisticated state-of-the-art algorithms using real-world ocean current data.

The rest of the paper is organized as follows. Sec. II briefly reviews some related literature. In Sec. III we detail the general trajectory planning problem formulation besides assumptions and conditions under which desired regret bounds hold. Sec. IV details the energy efficient trajectory planning problem formulation and solution methodology. Experimental evaluations are carried out to validate the performance of the proposed method, and the results are discussed in Sec. V. Finally, Sec. VI concludes the paper.

II. Related Work

Here we present a comparison of our work with existing literature as per the following criteria:

A. Ocean trajectory design approaches

The work done on trajectory planning in the ocean environments can be divided into graph-search based approaches which use variants of the A* algorithm, sampling based, and optimization techniques.

Graph-based search methods are designed to leverage the coarse resolution of data available [13]–[16]. However, traditional approaches suffer from two general problems: (i) the quality and complexity of the solution is controlled by the chosen level of discretization and (ii) the predicted trajectories often result in increased control costs because of the need to make sharp turns.

Sampling-based methods represent the trajectory as a sample from a Gaussian process [17], or sample a set of noisy trajectories about the initialized trajectory [12], [4]. Of these, [4] also samples in time, thereby providing additional flexibility in the design of trajectories. An exploration parameter is often required to control the variance of the sampled trajectories and a higher exploration parameter often leads to better but possibly more skewed trajectories. On the flip side, such sampling-based algorithms are highly sensitive and output a different trajectory for every run of the algorithm. Approaches such as [4] and [12] also require dense sampling rendering them slow in the presence of multiple constraints. Overall, the per-iteration run-time of sampling-based algorithms is at least $\mathcal{O}(KN^3)$ where K is the number of samples and N is the total number of waypoints. In contrast the proposed algorithm has only a single tuning parameter, no random components, and the time complexity of $\mathcal{O}(N)$.

Optimization methods approach the trajectory design problem more directly. Offline approaches include methods utilizing parallel swarm search [7], level-set expansion methods [18], [19], and gradient-based approaches [20]. These approaches rely on forecast data and cannot be readily adapted to online and time-varying settings. Online approaches to trajectory design include the sequential convex optimization framework in [21] and covariant gradient descent algorithm [22]. However, existing online approaches require static settings and very little analysis exists for the time-varying setting at hand.

Of these, most works provide time-optimal trajectory designs [7], [18]–[20]. In contrast, the final trajectory of the sampling-based algorithms such as [4] may be different every time depending on the initialization, sampled trajectories and ocean current velocities. Finally, two separate optimization problems, one for minimizing the travel time, and another for minimizing the energy, are posed in [16]. Although [16] is demonstrated under time-invariant flow fields, it can be easily extended to time-varying scenarios. However, extensions to time-varying flow fields limit the accuracy of the flow model used to predict future temporal variations. In contrast, the current work does not require such models and designs trajectories in an online fashion.

B. Environmental dynamics

Environmental dynamics considered here are limited to ocean currents. Papers such as [20] discuss planning in an estuarine environment in which the currents can be bi-directional, larger in magnitude than the vehicle's maximum velocity, and also temporally vary at the same location. Strong temporally varying currents are also discussed in [15] and [7]. While designing trajectories is significantly more challenging in such environments, none of the existing algorithms is capable of running in an online fashion.

To the best of our knowledge, the present work provides the first online utility-optimal trajectory design algorithm for application in time-varying oceanic environments and moving goal scenarios. The proposed OGD algorithm builds upon [2], where the OGD algorithm was first applied to the trajectory design problem. Different from [2], the present work provides a more generic variant of the OGD algorithm by considering time-varying coupling constraints in the optimization problem and its application to planning in oceanic environments.

III. General Trajectory Optimization Problem

This section formulates the general utility-optimal online trajectory learning problem and puts forth an online gradient descent (OGD) algorithm for the same. Considering a generic utility function and path constraints, we begin with formulating a constrained optimization problem that would serve as a benchmark for the proposed algorithm. Subsequently, the OGD formalism is used to develop the online algorithm with a sublinear offline regret. The algorithm and the formulation in this section build upon the trajectory design problem considered in [2] in the context of communication networks. The present formulation considers more general time-varying constraints and is therefore applicable to a broader variety of trajectory design problems.

The trajectory learning problem is expressed as that of finding the optimal agent locations $\mathbf{x}(t) \in \mathbb{R}^2$ at each $t \in \mathbb{N}$ subject to various restrictions. In the literature, these locations are also referred to as waypoints that the agent must traverse and the trajectory is simply the collection $\{\mathbf{x}(t)\}_{t\geq 1}$.

At time $t=1$, the current location of the agent $\mathbf{x}(1)=\mathbf{s}$, as well as the goal location $\mathbf{d}(1)$ are given. The agent learns the trajectory in an online fashion while the goal location $\mathbf{d}(t)$ continues to vary with t. Let T_{ETA} denote the expected time of arrival (ETA) for the agent when it follows the straight line path between $\mathbf{x}(1)$ and $\mathbf{d}(1)$ calculated either using forecast data or the prevailing environmental conditions at time $t=1$. However, due to the dynamic nature of the problem, at time $t=T_{\text{ETA}}$, (a) the agent need not be at $\mathbf{d}(0)$ if it follows the straight line path; and (b) the goal location would have already changed to $\mathbf{d}(T_{\text{ETA}})$. Besides, a budget of δ excess time slots is available and may be used to traverse a more energy-optimal trajectory. To this end, the general utility-optimal trajectory design problem may be formulated as follows

$$\{\mathbf{x}^\star(t)\}_{t=1}^T = \underset{\{\mathbf{x}(t)\}_{t=1}^T}{\arg\max} \sum_{t=1}^T U_t(\mathbf{x}(t)) \tag{1a}$$

$$\text{s. t.} \quad g_t(\mathbf{x}(t),\mathbf{x}(t+1)) \leq 0 \qquad 1 \leq t \leq T-1 \tag{1b}$$

$$\mathbf{x}(1) = \mathbf{s} \tag{1c}$$

$$\mathbf{x}(t) \in \mathcal{X} \qquad 1 \leq t \leq T \tag{1d}$$

where $T = T_{\text{ETA}} + \delta$ with δ representing excess time slots, $\mathcal{X} \subset \mathbb{R}^2$ is a non-empty closed and convex set that represents the functional space in which the agent moves, $U_t : \mathcal{X} \to \mathbb{R}$ is a time-varying concave utility function, and $g_t : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ represents the convex coupling constraint between successive way-points. Note that we denote the cost function $f_t := -U_t$ so as to express (1) as a more standard minimization problem. Further, both f_t and g_t may depend on other physical parameters of the agent/environment such as the goal location $\mathbf{d}(t)$, maximum agent speed $v_{\max}$ (measured in meters per time-slot), velocity of the ocean currents etc. Here, $\mathbf{x}^\star(1)=\mathbf{s}$ is the starting point and is therefore not an optimization variable in (1). While the subsequent analysis will require that the goal trajectory $\{\mathbf{d}(t)\}_{t=1}^T$ and the functions f_t and g_t to be slowly time-varying, their temporal variations are otherwise arbitrary and possibly adversarial.

Problem (1) may be solved in an offline fashion using any existing convex optimization algorithm. In practice, however, an online solution is desirable since the state of the environment at time t (encoded in the functions f_t and g_t) is revealed after the action at time t is taken. The setting is reminiscent of the online convex optimization (OCO) framework which however does not include temporally coupled constraints (1b). Towards solving (1) in an online manner using online gradient descent (OGD) algorithm that takes the following form for $t \geq 1$

$$\hat{\mathbf{x}}(t+1) = \hat{\mathbf{x}}(t) - \frac{1}{\gamma_t}\nabla f_t(\hat{\mathbf{x}}(t)), \tag{2}$$

where γ_t is the step-size and the algorithm is initialized with $\hat{\mathbf{x}}(1)=\mathbf{s}$. Although OGD has been widely applied to unconstrained OCO problems, its performance for temporally constrained problems has not been well-studied. A special case of (1) was considered in [2] for the case when $g_t(\mathbf{x},\mathbf{x}') = \|\mathbf{x}-\mathbf{x}'\| - v_{\max}$ for some $v_{\max} \in \mathbb{R}_{++}$. The current work considers a more general setting with arbitrary convex coupling function g_t.

We begin by stating some regularity conditions for the problem (1) that are necessary to develop meaningful guarantees.

- **(A1)** **Strong convexity:** the function f_t is μ-convex, i.e., $f_t(\mathbf{x}) - \frac{\mu}{2}\|\mathbf{x}\|^2$ is convex.
- **(A2)** **Lipschitz continuous gradient:** the function f_t is L-smooth, i.e., $\|\nabla f_t(\mathbf{x}) - \nabla f_t(\mathbf{x}')\| \leq L\|\mathbf{x}-\mathbf{x}'\|$ for all $\mathbf{x}$, $\mathbf{x}' \in \mathcal{X}$.
- **(A3)** **Bounded gradients**: the objective function gradients are bounded, i.e., $\|\nabla f_t(\mathbf{x})\| \leq G$ for all $\mathbf{x} \in \mathcal{X}$.
- **(A4)** **Feasibility:** the OGD iterates $\hat{\mathbf{x}}(t)$ adhere to (1b)-(1d).

Observe here that while Assumptions (A1)-(A3) are standard, Assumption (A4) is specific to the problem at hand. In general, it is required that Assumption (A4) be explicitly checked. In the present case, we choose the step size γ_t to ensure that (A4) continues to hold; see Sec. IV.

These set of assumptions enable us to characterize the performance of the OGD algorithm through the notion of regret, which compares the cost incurred by the proposed algorithm against that incurred by an adversary that has complete information about future changes and solves the problem in an offline manner. Here we use the following definition of regret that is motivated from the offline regret introduced in [23],

$$\mathbf{Reg}_T := \underbrace{\left[\sum_{t=2}^T f_t(\hat{\mathbf{x}}(t))\right]}_{\text{online}} - \underbrace{\left[\sum_{t=2}^T f_t(\mathbf{x}^\star(t))\right]}_{\text{offline}}. \tag{3}$$

The regret bound in (3) is calculated in terms of the squared path length of the adversary, defined as

$$S_T := \sum_{t=1}^{T-1} \|\mathbf{x}^\star(t+1) - \mathbf{x}^\star(t)\|^2, \tag{4}$$

and squared cumulative gradient variation as

$$\mathcal{G}_T := \sum_{t=1}^{T-1} \max_{\mathbf{x}\in\mathcal{X}} \|\nabla f_{t+1}(\mathbf{x}) - \nabla f_t(\mathbf{x})\|_2^2. \tag{5}$$

Intuitively, in the former case, significant variations in $\mathbf{x}^\star(t)$ make it difficult for the agent to follow and may lead to a linear regret. For instance, if the squared path length is linear in T, e.g., if the goal location $\mathbf{d}(t)$ moves away from the agent by $v_{\max}$ meters per time slot, the adversary may never be able to catch up with the goal. The cumulative gradient variation is also similar and has been widely used in the context of online learning and dynamic optimization [24], [25].

Let $\bar{\gamma} := \max_{1\leq t\leq T-1}\gamma_t$ and $\underline{\gamma} := \min_{1\leq t\leq T-1}\gamma_t$ be the constants that do not depend on T. The following theorem summarizes the regret bound for the OGD algorithm (2).

Theorem 1: Under the assumptions (A1)-(A4) and for $\mu < \bar{\gamma} < 2\underline{\gamma} - L$, the sequence of $\hat{\mathbf{x}}(t)$ generated by (2) adheres to the regret bound

$$\mathbf{Reg}_T \leq \mathcal{O}\sqrt{T(S_T + \mathcal{G}_T)}. \tag{6}$$

The result in *Theorem 1* states that for large values of T and for a sub-linearly time-varying adversary, the online algorithm incurs a sub-linear regret over the off-line solution. The proof of Theorem 1 follows along the lines of that in [2] but includes modifications required to handle the generic time-varying convex constraint function g_t. Proof to *Theorem 1* is made available in the supplementary material[1].

IV. Trajectory Planning in Ocean Environments

The proposed utility-optimal trajectory design algorithm is applied to a watercraft operating in an ocean environment and seeking a possibly time-varying goal. While the watercraft has propulsive capabilities, it is required to reach the goal in a limited amount of time while expending minimal energy. The following information is available at any time t:

- Historical data on ocean currents in the region such as that available from Regional Ocean Modeling System (ROMS) [9],
- Current watercraft location $\hat{\mathbf{x}}(t)$,
- Current ocean velocity $\mathbf{v}_o(t)$ at location $\hat{\mathbf{x}}(t)$, and
- Current goal location $\mathbf{d}(t)$,

in addition to problem parameters such as T, maximum watercraft velocity in still water $v_{\max}$, and other tunable parameters.

Designing online trajectories in such time-varying and uncertain environments is challenging and has never been attempted before. Existing methods usually rely on forecasts and require re-solving the full problem at every time slot. While such an approach can also be adopted here, we are interested in a more computationally efficient algorithm. The proposed algorithm will be based on the OGD algorithm in (2) and therefore adheres to the guarantees in Theorem 1. However, in order to ensure Assumption (A4) and to obtain reasonable performance, it is required to design the functions f_t and g_t as well as provide rules for choosing the step-size.

A. Design of the objective function f_t

We consider the objective function

$$f_t(\mathbf{x}(t)) := \lambda(t)\,\|\mathbf{x}(t)-\mathbf{d}(t)\|^2 + (1-\lambda(t))\langle \hat{\mathbf{x}}(t-1)-\mathbf{x}(t), \mathbf{v}_o(t)\rangle \tag{7}$$

where $\lambda(t) \in [0,1]$, is a control parameter. The objective function in (7) is, therefore, a convex combination of the current squared distance from the target $\|\mathbf{x}(t)-\mathbf{d}(t)\|^2$ and the component of vehicle velocity in the negative direction of the ocean velocity. Consequently, an appropriate choice of $\lambda(t)$ ensures that the planned trajectory is headed towards the current goal location and in the direction of ocean currents, if possible. For instance, when $\lambda(t)$ is close to 1, the watercraft ignores the ocean currents and heads straight to the goal. Clearly, the choice of $\lambda(t)$ is critical towards ensuring the goal is reached in a timely manner, and at the same time, energy consumption is minimized. Next, we discuss two different strategies for the choice of $\lambda(t)$.

[1]Proof of *Theorem 1* is provided at `https://goo.gl/3LLAvp`

TABLE I: Behavior of direction-dependent $\lambda(t)$ choice

Direction of ocean currents w.r.t goal	Magnitude of ocean currents		
	Strong $\eta(t)\approx 1$	Moderate $\eta(t)\approx 0.5$	Lower $\eta(t)\approx 0$
$\theta(t)\in[0,\frac{\pi}{2})$	$\lambda(t)\in[0,0.5)$	$\lambda(t)\in[0.5,0.75)$	$\lambda(t)\approx 1$
$\theta(t)\in[\frac{\pi}{2},\pi]$	$\lambda(t)\in[0.5,1]$	$\lambda(t)\in[0.75,1)$	$\lambda(t)\approx 1$

1) Increasing $\lambda(t)$: The choice $\lambda(t)=t/T$ for $1\le t\le T$ is motivated from the observation that the watercraft may drift along the direction of the ocean currents initially to save energy. However, as time goes on, more importance must be placed on reaching the goal. While such a strategy was used in [2], it is agnostic to the direction and magnitude of $\mathbf{v}_o(t)$ and therefore suboptimal (see Fig.1(b)). In particular, this strategy fails when the ocean currents are pointing directly away from the goal, ultimately resulting in more energy expenditure than for the straight line path. Note that with this choice of $\lambda(t)$, Assumptions (A1)-(A3) are satisfied.

2) Direction-dependent $\lambda(t)$: Instead of *increasing* $\lambda(t)$, information about the direction of ocean currents may be taken into account using the following choice:

$$\lambda(t) = 1-\eta(t)\cos^2\left(\frac{\theta(t)}{2}\right) \tag{8}$$

where $\theta(t) = \angle(\mathbf{d}(t)-\hat{\mathbf{x}}(t), \mathbf{v}_o)$, $\eta(t) = \frac{\|\mathbf{v}_o(t)\|}{v_o^{\max}}$, and $v_o^{\max}$ is the maximum ocean currents velocity estimated from historical data. Here, the ratio $\frac{\|\mathbf{v}_o(t)\|}{v_o^{\max}}$ is the relative strength of the ocean currents at a given point. When the ocean currents are weak, i.e., $\|\mathbf{v}_o(t)\| \ll v_o^{\max}$ or when the currents are directed away from the goal location, i.e., $\theta\approx\pi$ radians, $\lambda(t)$ is close to 1 and consequently the watercraft is headed towards the target. However, when the ocean currents are in a favorable direction $\lambda<1$ and therefore the watercraft may utilize the ocean velocity in order to save energy. Table I provides some example values of $\lambda(t)$ for various values of $\eta(t)$ and $\theta(t)$.

B. Design of the constraint function g_t

Recalling that the average velocity of the watercraft at time slot t is $\mathbf{x}(t+1)-\mathbf{x}(t)$ and measured in meters per slot, we consider the constraint function of the form

$$g_t(\mathbf{x}(t+1),\mathbf{x}(t)) := \|\mathbf{x}(t+1)-\mathbf{x}(t)-\mathbf{v}_o(t)\| - \alpha(t)v_{\max} \tag{9}$$

where $\mathbf{x}(t+1)-\mathbf{x}(t)-\mathbf{v}_o(t)$ represents the velocity of the watercraft relative to the ocean. Note that the magnitude of the relative velocity of the watercraft is physically limited to at most $v_{\max}$ meters per slot. The additional factor of $\alpha(t)\in(0,1]$ is included to further restrict the maximum relative velocity if required. When the objective function with $\lambda(t)$ as in (8) is used, it only provides directional information and $\alpha(t)$ must be carefully chosen to ensure that Assumption (A4) is satisfied. Towards this end, we consider

$$\alpha(t) = \exp\left(-\beta\left(\frac{\delta}{T_{\mathrm{ETA}}}+\eta(t)\cos\left(\frac{\theta(t)}{2}\right)\right)\right) \tag{10}$$

where δ, $\eta(t)$, and $\theta(t)$ are as defined earlier. Intuitively, $\alpha(t)$ is small when the ocean currents are in a favorable direction ($\theta\ll\pi/2$) and are either strong (large $\eta(t)$) or

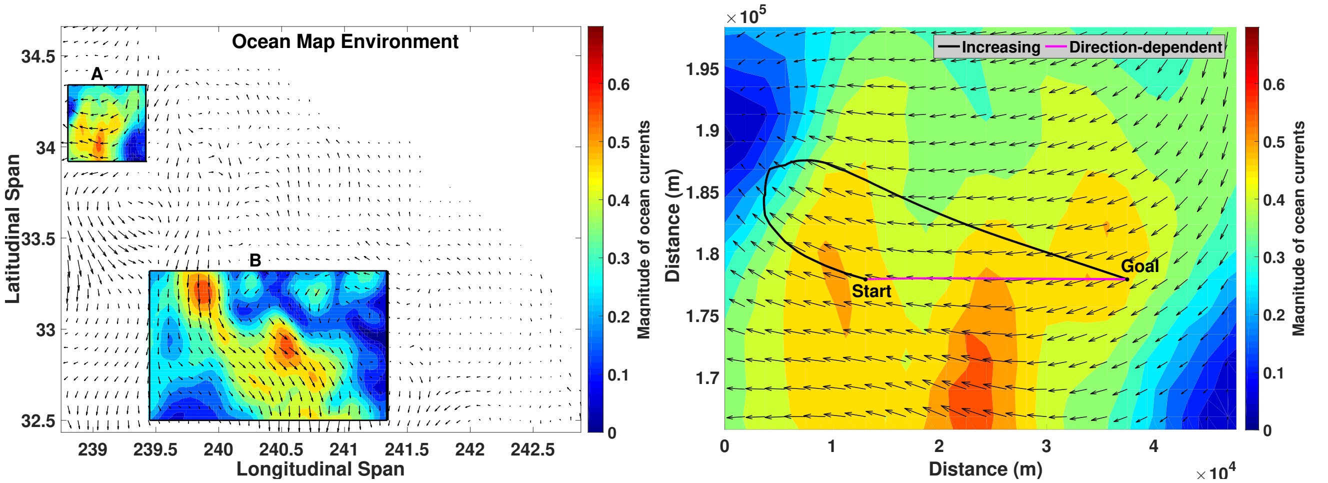

(a) (b)

Fig. 1: (a) Southern California Bight: Note that the ocean currents magnitudes at islands are made negligible (b) Demonstrating the behavior of Increasing vs. Direction-dependent strategies for $\lambda(t)$.

sufficient excess time is available (large δ) to reach the goal. In such a scenario, the watercraft simply drift along the currents and can conserve energy. On the other hand, the watercraft runs the engine at full capacity when $\alpha(t)$ is close to one, such as when the ocean currents are in the opposite direction. The tuning parameter β is dependent on the environment and must be learned a priori using historical data.

Finally, the OGD updates are applied using the objective function in (7) while the step-size γ_t is chosen to satisfy Assumption (A4) where g_t is chosen as per (9). We refer to our algorithm as online trajectory optimization using online gradient descent (OTOGD).

V. Experimental Evaluations

We tested our algorithm in a real oceanic environment built from historical ROMS ocean currents data taken from [9]. Specifically, we have chosen Southern California Bight output collected on August 10, 2018, as shown in Fig. 1(a). From the dataset, it is found that the ocean currents' velocity is found to be varying from 0.001 m/s to a maximum of 0.69 m/s. Thereby we assume that the maximum velocity of the watercraft in still water is $1 m/s$ to make sure that it can travel forward even in strong disturbances. The energy incurred for travelling in planned trajectories is calculated in the same manner as detailed in [7] and is given by

$$E = c_d V_r^3 t \tag{11}$$

where c_d is the vehicles' drag coefficient, V_r is the magnitude of the required velocity for the motors to provide and t is the travel time for the relevant section of the path.

A. Static goal

We first illustrate the working of the algorithm when the goal is static and compare the different strategies for choosing $\lambda(t)$. A chunk of the real data labelled 'A' (see Fig. 1(a)) is considered where the ocean currents with high magnitude are directed away from the goal. The start and goal locations are as specified in Fig. 1(b). It can be observed that the OTOGD algorithm with *increasing* $\lambda(t)$ makes the vehicle go away from the goal initially when $\lambda(t)$ is small, and as time goes on, $\lambda(t)$ increases and more emphasis is placed on reaching the goal. Therefore the vehicle steadily starts going in the direction of the goal. Such a trajectory is clearly suboptimal since it requires the vehicle to travel a longer distance against the current. In contrast, the *direction-dependent* choice of $\lambda(t)$ makes the vehicle travel directly to the goal and incurs less energy, as also evident from Table II.

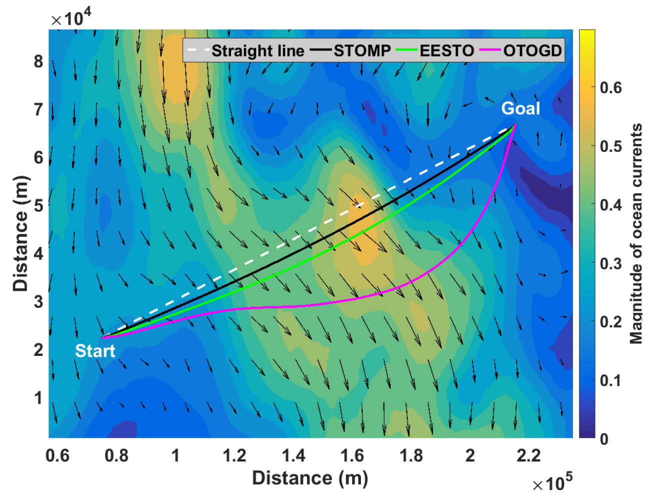

Fig. 2: Trajectory comparisons of STMOP [12], EESTO [4] and OGD with energy efficient control strategy. Note that straight line trajectory with broken lines is shown for reference.l

TABLE II: Energy cost comparisons

$\lambda(t)$	**Energy cost (kJ)**
Increasing	184.42
Direction-dependent	130.29

To illustrate the performance of the proposed algorithm in comparison to the state-of-the-art algorithms under a static

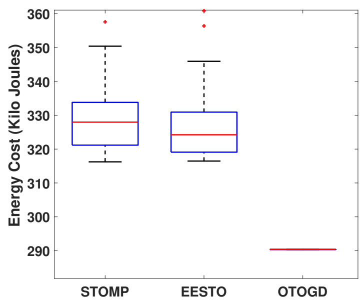

Fig. 3: Energy cost comparisons for STOMP [12], EESTO [4], OTOGD

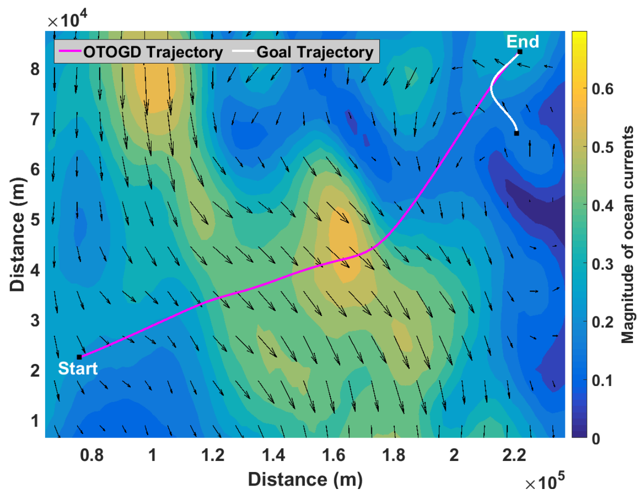

Fig. 4: Illustration of planning in moving goal scenario

goal scenario, we now choose another chunk of real data 'B' as shown in Fig. 1(a). The start and goal locations are as specified in Fig. 2 and are separated by around 147 km. Trajectories are generated using STOMP [12], EESTO [4] and the proposed OTOGD with the *direction-dependent* control strategy. Fig. 2 displays those trajectories along with the ocean current disturbances. Since STOMP [12] and EESTO [4] are sampling based methods they do not output the same trajectory at every run. Instead, we executed the algorithms 100 times each and picked the trajectory with the minimum energy cost. The total energy consumed for all three trajectories is also shown in Fig. Fig. 3. It is evident from Fig. 3 that the proposed algorithm is generating trajectories that are more energy efficient than the state-of-the-art methods besides being online. More importantly, unlike STOMP [12] and EESTO [4], the obtained trajectory is the same for each run, and the resulting output has zero variance.

B. Dynamic goal

To illustrate the performance of the proposed algorithm under dynamic goal scenario, we make use of the same chunk 'B' of the real data keeping the start location the same as earlier. Here we simulated a moving goal scenario and assumed that the exact location of the goal at every time instant is known to the vehicle. Note that the proposed framework still requires the maximum velocity of the goal to

TABLE III: Energy conservation vs. delay analysis

% of excess time slots	Energy conserved (kJ)		
	EESTO	STOMP	OTOGD
10%	112.01 ± 24.4	107.10 ± 28.2	123.08
30%	184 ±26.6	170.2 ± 31.1	217.13

be significantly less than that of the vehicle. Fig. 4 displays the final trajectory along with the trajectory traversed by the goal. Observe the change in the shape of the trajectory as compared to that of the static goal scenario in Fig. 2.

It is remarked that in theory, existing algorithms such as [4] can perform re-planning at every time instant thereby following a moving goal. However, with per-waypoint re-planning, EESTO incurs a computational cost of $\mathcal{O}(IKN^4)$ where K is the number of samples, N represents the total number of waypoints, and I is the total number of iterations required for convergence. In contrast, the proposed algorithm only incurs a cost of $\mathcal{O}(N)$ even when the goal and the environment are time-varying.

C. Delay vs energy conservation analysis

Here we ran STOMP [12], EESTO [4] along with the proposed OTOGD algorithm for 10% and 30% allowed delay in time for 30 randomly chosen chunks of real data displayed in Fig 1(a). For each of the case, Table III shows the average amount of energy conserved as compared to the energy incurred in travelling along the straight line path from the starting point to the goal. As evident from the table, the proposed algorithm saves more energy as compared to state-of-the-art algorithms. Further, existing algorithms exhibit a high variance in the conserved energy, requiring multiple runs.

VI. Conclusions

This paper considered the problem of designing an online framework for energy-efficient trajectory planning under strong and time-varying disturbances in oceanic environments. The problem is formulated as a time-varying constrained convex optimization problem and solved using online gradient descent updates adhering to motion constraints. Leveraging the online framework, the proposed algorithm is capable of handling temporal variations in both the ocean currents as well as the goal locations and incurs sub-linear offline regret. Numerical tests carried out using historical regional ocean modelling system dataset, establish that the proposed algorithm yields energy efficient trajectories as compared to the state-of-the-art algorithms. Besides, the proposed algorithm is computationally efficient by several orders of magnitude. Future work includes incorporating static or dynamic obstacle avoidance techniques and handling uncertainties in the vehicle or goal locations.

VII. ACKNOWLEDGMENT

The authors would like to thank Lavish Arora (email: lavi@iitk.ac.in), Bharadwaj Jampu (email: jampub@iitk.ac.in) and Amrit S. Bedi (email: amrit0714@gmail.com) for their participation in the discussions on trajectory optimization in oceanic environments.

REFERENCES

[1] M. Hoy, A. S. Matveev, and A. V. Savkin, "Algorithms for collision-free navigation of mobile robots in complex cluttered environments: a survey," *Robotica*, vol. 33, no. 3, pp. 463–497, 2015.

[2] A. S. Bedi, K. Rajawat, and M. Coupechoux, "An online approach to D2D trajectory utility maximization problem," in *Proc. IEEE Conference on Computer Communications (INFOCOM)*, April 2018.

[3] M. Cirillo, F. Pecora, H. Andreasson, T. Uras, and S. Koenig, "Integrated motion planning and coordination for industrial vehicles," in *Proc. International Conference on Automated Planning and Scheduling (ICAPS)*, 2014.

[4] D. Jones and G. A. Hollinger, "Planning energy-efficient trajectories in strong disturbances," *Proc. IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2088–2095, 2017.

[5] R. N. Smith, Y. Chao, P. P. Li, D. A. Caron, B. H. Jones, and G. S. Sukhatme, "Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model," *Proc. International Journal of Robotics Research*, vol. 29, no. 12, pp. 1475–1497, 2010.

[6] D. Verscheure, B. Demeulenaere, J. Swevers, J. De Schutter, and M. Diehl, "Time-energy optimal path tracking for robots: a numerically efficient optimization approach," in *Proc. International Workshop on Advanced Motion Control*, 2008, pp. 727–732.

[7] J. Witt and M. Dunbabin, "Go with the flow: Optimal AUV path planning in coastal environments," in *Australian Conference on Robotics and Automation (ACRA)*, vol. 2008, no. 2, 2008.

[8] B. Garau, A. Alvarez, and G. Oliver, "Path planning of autonomous underwater vehicles in current fields with complex spatial variability: an a* approach," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2005, pp. 194–198.

[9] A. F. Shchepetkin and J. C. McWilliams, "The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model," *Ocean modelling*, vol. 9, no. 4, pp. 347–404, 2005.

[10] P. J. Martin, "Description of the navy coastal ocean model version 1.0," Naval Research Lab STENNIS SPACE CENTER MS, Tech. Rep., 2000.

[11] G. A. Hollinger, A. A. Pereira, J. Binney, T. Somers, and G. S. Sukhatme, "Learning uncertainty in ocean current predictions for safe and reliable navigation of underwater vehicles," *Journal of Field Robotics*, vol. 33, no. 1, pp. 47–66, 2016.

[12] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "STOMP: Stochastic trajectory optimization for motion planning," *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4569–4574, 2011.

[13] P. Svec, M. Schwartz, A. Thakur, and S. K. Gupta, "Trajectory planning with look-ahead for unmanned sea surface vehicles to handle environmental disturbances," in *Proc. IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2011, pp. 1154–1159.

[14] T. Lee, H. Kim, H. Chung, Y. Bang, and H. Myung, "Energy efficient path planning for a marine surface vehicle considering heading angle," *Ocean Engineering*, vol. 107, pp. 118–131, 2015.

[15] V. T. Huynh, M. Dunbabin, and R. N. Smith, "Predictive motion planning for AUVs subject to strong time-varying currents and forecasting uncertainties," in *Proc. IEEE International Conference on Robotics and Automation (ICRA),*, 2015, pp. 1144–1151.

[16] D. Kularatne, S. Bhattacharya, and M. A. Hsieh, "Time and energy optimal path planning in general flows." in *Robotics: Science and Systems*, 2016.

[17] M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots, "Continuous-time gaussian process motion planning via probabilistic inference," *The International Journal of Robotics Research*, vol. 37, no. 11, pp. 1319–1340, 2018.

[18] D. N. Subramani and P. F. Lermusiaux, "Energy-optimal path planning by stochastic dynamically orthogonal level-set optimization," *Ocean Modelling*, vol. 100, pp. 57–77, 2016.

[19] T. Lolla, P. F. Lermusiaux, M. P. Ueckermann, and P. J. Haley, "Time-optimal path planning in dynamic flows using level set equations: theory and schemes," *Ocean Dynamics*, vol. 64, no. 10, pp. 1373–1397, 2014.

[20] D. Kruger, R. Stolkin, A. Blum, and J. Briganti, "Optimal AUV path planning for extended missions in complex, fast-flowing estuarine environments," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2007, pp. 4265–4270.

[21] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.

[22] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "CHOMP: Gradient optimization techniques for efficient motion planning," *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, pp. 489–494, 2009.

[23] T. Chen, Q. Ling, and G. B. Giannakis, "An online convex optimization approach to proactive network resource allocation," *IEEE Transactions on Signal Processing*, vol. 65, no. 24, pp. 6350–6364, 2017.

[24] L. Zhang, T. Yang, J. Yi, J. Rong, and Z.-H. Zhou, "Improved dynamic regret for non-degenerate functions," in *Advances in Neural Information Processing Systems (NIPS)*, 2017.

[25] O. Besbes, Y. Gur, and A. Zeevi, "Non-stationary stochastic optimization," *Operations research*, vol. 63, no. 5, 2015.