



Risk-optimal path planning in stochastic dynamic environments

Deepak N. Subramani, Pierre F.J. Lermusiaux*

Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Rm 5-207B, Cambridge, MA 02139, United States

Received 12 January 2019; received in revised form 21 April 2019; accepted 23 April 2019
Available online xxxx

Highlights

- Combined decision theory and time-optimal S-PDEs in uncertain, dynamic, strong flows.
- Risk of being suboptimal given the uncertain time-optimal path predictions minimized.
- Multiple error and cost metrics used for rigorous risk evaluation and minimization.
- Applied planning in stochastic front, double-gyre QG flow and flow exiting a strait.

Abstract

We combine decision theory with fundamental stochastic time-optimal path planning to develop partial-differential-equations-based schemes for risk-optimal path planning in uncertain, strong and dynamic flows. The path planning proceeds in three steps: (i) predict the probability distribution of environmental flows, (ii) compute the distribution of exact time-optimal paths for the above flow distribution by solving stochastic dynamically orthogonal level set equations, and (iii) compute the risk of being suboptimal given the uncertain time-optimal path predictions and determine the plan that minimizes the risk. We showcase our theory and schemes by planning risk-optimal paths of unmanned and/or autonomous vehicles in illustrative idealized canonical flow scenarios commonly encountered in the coastal oceans and urban environments. The step-by-step procedure for computing the risk-optimal paths is presented and the key properties of the risk-optimal paths are analyzed.

© 2019 Elsevier B.V. All rights reserved.

Keywords: Stochastic path planning; Level set equations; Dynamically orthogonal; Ocean modeling; AUV; Uncertainty quantification

1. Introduction

In recent years, the use of autonomous platforms such as Autonomous Underwater Vehicles, gliders, floats, drones, and Unmanned Aerial Vehicles has rapidly increased, for both underwater and aerial applications. These vehicles operate in dynamic environments with strong and uncertain currents/winds that affect their motion. To pilot these vehicles efficiently between any two locations, it is important to utilize environmental predictions to plan paths that optimize travel time, energy consumption, data collection and/or safety. Such environmental predictions are, however, often uncertain and path planning has to rigorously account for these probabilistic predictions, if available. Recently, we developed governing stochastic partial differential equations and their efficient stochastic Dynamically

* Corresponding author.

E-mail addresses: deepakns@mit.edu (D.N. Subramani), pierrel@mit.edu (P.F.J. Lermusiaux).

Orthogonal (DO) counterparts for time-optimal path planning in uncertain, strong and dynamic flow fields [1]. With these equations, we can forecast the probability distribution of reachability fronts and time-optimal paths. Such probabilistic predictions raise important new questions: What is the optimal path choice under such uncertainty?; Can/should concepts from rational decision making be utilized for informing this optimal choice?; What is an appropriate risk measure and optimality criterion for such choices?; Can such paths be computed efficiently? We answer these questions in the present paper.

The present work is based on fundamental differential equations that govern the evolution of the reachability and time-optimal paths in strong and dynamic currents [2], and is related to energy-optimal paths [3,4]. These equations were used to compute optimal paths, both in realistic data-driven simulations [5,6] and with real vehicles [7]. They were also employed to solve pursuit evasion problems [8,9]. The input probabilistic predictions of flow fields are obtained from the variance-optimal reduced-order stochastic dynamically orthogonal equations [10–12]. The resulting stochastic PDEs (S-PDEs) for time-optimal path planning have some advantages [1]: (i) for a given stochastic environmental flow prediction, the stochastic time-optimal paths are exact; (ii) the computed paths naturally avoid stationary and dynamic obstacles; and (iii) the probability of a location being reachable (or non-reachable) is directly predicted.

From the probability distribution of time-optimal paths, we have to assess the risks and make a decision of risk-optimal paths. This subject of decision making under uncertainty has been well studied in the fields of economics and management. One widely used model is the expected utility theory and its several variants as reviewed in [13]. The key ingredients of this model are the evaluation of the utility cost of the outcome due to a decision and the probability of that outcome. The expected utility theory can be utilized in a prescriptive or normative framework to inform optimum decision making under complex decision scenarios and can be customized to fit the risk behavior of users [13–15]. Specific utility functions are also available for different risk behaviors [e.g., 16–18].

Our goal here is to combine a principled risk optimality criterion grounded in decision theory with our stochastic dynamically orthogonal level-set equations to develop efficient computational schemes to predict risk-optimal paths from the distribution of stochastic time-optimal paths. We also seek to apply the new schemes to several stochastic flow scenarios, analyze the effects of different risk metrics and criteria, and discuss the properties of risk-optimal paths.

The remainder of the present paper is organized as follows. Next, we provide a brief review of prior path planning results. In Section 2, we develop the theory and schemes for risk-optimal path planning. In Section 3, we apply the new schemes to compute risk optimal paths for a variety of stochastic flow scenarios. In Section 4, we conclude and provide future research directions.

1.1. Previous progress in optimal path planning

Traditionally, path planning has focused on land-based robots in stationary environments [e.g., 19–21]; however, a major challenge for marine and aerial platforms is the effects of the uncertain, strong and dynamic currents/winds. Several authors have extended many of the algorithms for static environments to plan paths of autonomous vehicles in dynamic environments (for reviews, see e.g., [2] and [22]). These applications include graph based search methods such as the modified Dijkstra's algorithm [23], A*-search [24], RRTs [25], kinematic tree-based navigation [26], stochastic planners with uncertain edge weights [27], and stochastic surface response methods [28]. Other techniques such as nonlinear optimization methods [29,30], sequential quadratic programming [31], evolutionary algorithms [32,33], fast marching methods [34,35], wave front expansion [36,37] have been used with varying degrees of success. Monte Carlo methods to account for uncertainties and compute statistics of optimal trajectories have been used with potential field methods [38] and two point boundary value problems [39]. Partially Observable Markov Decision Processes (POMDP) have been utilized for computing policies of robots moving in uncertain environments [40]. MDP and graph search methods have also been recently utilized for accounting for forecasting uncertainties while planning [41]. We refer to [42,43] and references therein for a detailed review of all these results, and to Blythe [44] for an early review and outlook of decision-theoretic planning.

Previous attempts to combine expected utility theory with robot motion planning in dynamic environments have been limited. Rapidly-exploring Random Trees (RRTs) with utility-based random trees have been used for single-query robot planners [45]. Here, the authors leverage the solution of a local linear dynamics, quadratic cost, and Gaussian noise (LQG) control problem to predict a path distribution from which a path that minimizes

the expected cost is computed through incremental sampling refinement. In Boddy and Dean [46], the authors propose an expectation driven iterative refinement approach with a heuristic, but without providing any algorithms or applications. In Burns and Brock [47], the expected utility of a path to be collision free is utilized with A* search to iteratively identify paths that maximize utility. The authors estimate the probability of an obstacle free path by noisy sensing and demonstrate the planning algorithm with a simple laboratory experiment. A maximum expected-utility-based central arbiter for the Distributed Architecture for Mobile Navigation (DAMN) framework has been developed and showcased for land robots [48]. The commands for future motion under DAMN are then decided by choosing the trajectory that maximizes the expected utility of steering the robot through obstacle free regions. In Pereira et al. [22], a risk-aware planner based on a minimum expected risk criterion and using a Markov Decision Process has been demonstrated to plan collision free paths for gliders in simulations and field experiments. The planning was then aware of currents, but did not actively use them to optimize travel parameters. For Unmanned Aerial Vehicles (UAV), De Filippis et al. [49] have developed an algorithm that utilizes collision risk maps and graph search algorithm to plan paths to minimize such risk. In Rudnick-Cohen et al. [50], the planning algorithm minimizes the risk to persons on the ground when a UAV flies over inhabited areas. Another focus of researchers has been to minimize the risk of capture of autonomous vehicles moving in environments with static and known threats, for example optimal paths that avoid radars by a graph search scheme [51], or by solving Boundary Value Problems [52].

Briefly, the main issue with several of the above methods is that they are inaccurate, or apply specific heuristics, or are computationally intractable in strong, dynamic, and uncertain flows. Also, several methods do not rigorously utilize environmental predictions for planning. Additionally, accurate accounting of uncertainties requires large ensemble sizes for convergence of Monte Carlo methods, making them computationally very expensive for real-time use. The focus of previous works utilizing expected utility theory has been mostly on minimizing the risk of collision within uncertain environments. They have not yet been utilized for rigorous risk-optimal planning with PDEs in stochastic dynamic environments.

2. Theory and schemes

2.1. Problem statement

The risk-optimal path planning problem in uncertain flows can be formulated as follows (Fig. 1). Consider a domain \mathcal{D} with a spatial index \mathbf{x} , temporal index t , probabilistic sample space Ω , random event $\omega \in \Omega$, and probability distribution function $p_{\Omega}(\bullet)$. For a vehicle P navigating from start \mathbf{x}_s to a target \mathbf{x}_f in a stochastic flow $\mathbf{v}(\mathbf{x}, t; \bullet)$, we denote the time-optimal path distribution by $\mathbf{X}_P(\mathbf{x}_s, t; \bullet)$ and the time-optimal heading distribution by $\hat{h}(t; \bullet)$. We define risk-optimal paths as the time-optimal paths that minimize the expected utility cost of following a path that is possibly sub-optimal in the uncertain environment. Next, we develop the theory to compute such risk-optimal paths.

2.2. Theory of risk-optimal paths

Risk-optimal paths can be obtained in three steps. First, $\mathbf{v}(\mathbf{x}, t; \bullet)$ is computed by solving the DO stochastic environmental flow equations (Appendix A.1; [11,53]). Second, $\mathbf{X}_P(\mathbf{x}_s, t; \bullet)$ and $\hat{h}(t; \bullet)$ are computed by solving the stochastic dynamically orthogonal level-set Hamilton Jacobi equations (A.7) followed by the stochastic backtracking equation (A.5). We developed and applied these equations in [1]. Third, a risk metric is selected, the risk of all path choices $\mathbf{X}_P(\mathbf{x}_s, t; \bullet)$ is evaluated, and the risk-optimal path $\mathbf{X}_P^*(\mathbf{x}_s, t)$ computed by minimizing the risk. Here, we focus on all aspects of this third step.

To evaluate and minimize the risk of following a wrong path from among the distribution of time-optimal paths, we employ the expected utility hypothesis [e.g.,13,16]. The risk corresponding to a path choice is the expected utility cost of that choice being sub-optimal due to the uncertain environment. The utility cost depends on the error incurred due to the sub-optimality and the cost a decision maker assigns to this error, in accordance with their risk profile or chosen risk metric. Mathematically, we can formulate the risk evaluation and minimization problem as follows.

The time-optimal path $\mathbf{X}_P(\mathbf{x}_s, t; \omega_l)$ or heading time-series $\hat{h}(t; \omega_l)$ corresponding to a random event $\omega_l \in \Omega$ is exact for the environmental realization $\mathbf{v}(\mathbf{x}, t; \omega_l)$; however, its use in another environmental realization $\mathbf{v}(\mathbf{x}, t; \omega_m)$

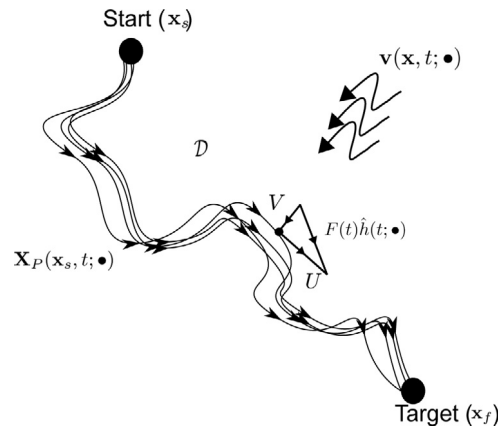


Fig. 1. Schematic of minimum-risk time-optimal path planning setup: The goal is to compute the time-optimal path with minimum risk under uncertainty for a vehicle P navigating from \mathbf{x}_s to \mathbf{x}_f in an uncertain flow field $\mathbf{v}(\mathbf{x}, t; \omega)$. The effective velocity, U experienced by the vehicle is the vector sum of the vehicle’s forward motion $F(t)\hat{h}(t)$ and the background flow V .

would be sub-optimal, and potentially infeasible. Let the trajectory achieved by following the waypoints or headings of the optimal path for ω_l in ω_m be $\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m)$. We define the error incurred due to this path as $e(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m))$ and the cost $C(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m)) = f(e(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m)))$ to quantify the effect of this sub-optimality on the decision maker. Here, e is an error function selected by the operator as the error metric of choice and f is a utility cost function corresponding to the risk profile of the decision maker. These functions e and f are discussed and defined in Section 2.3. Thus, the risk of utilizing the time-optimal path $\mathbf{X}_P(\mathbf{x}_s, t; \omega_l)$ is the expected cost, i.e,

$$R(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l)) = \int_{\Omega} C(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m))P_{\Omega}(\mathbf{v}(\mathbf{x}, t; \omega_m))d\omega_m. \tag{1}$$

Then, the risk-optimal path is computed by optimization as

$$\mathbf{X}_P^*(\mathbf{x}_s, t) = \arg \min_{\mathbf{X}_P(\mathbf{x}_s, t; \omega_l)} R(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l)). \tag{2}$$

Of course, there can be many time-optimal paths that have the same risk $R(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l))$ in Eq. (1). Similarly, the risk-optimal path $\mathbf{X}_P^*(\mathbf{x}_s, t)$ defined in Eq. (2) can also be multi-valued.

2.3. Schemes for risk evaluation and minimization

The specific choice of the error incurred due to sub-optimality e and the form of utility cost function f determine the risk evaluation and risk-optimal path selection. For brevity of notation, we use hereafter e_{lm} for the error $e(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m))$ and C_{lm} for the cost $C(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m))$. There can be several error metrics e corresponding to the operational parameters of interest and multiple analytical forms of the utility cost function f corresponding to different risk tolerance profiles of different decision makers, e.g., the pilot or mission designer [13,18]. Next, we present the schemes we consider.

2.3.1. Error metric e_{lm}

The purpose of the error metric in our problem is to quantify the difference between following a prescribed path (specified as waypoints or a heading time-series) in an environment for which it is sub-optimal and following the true optimal path for that environment. We consider vehicle operations where the objectives are either achieving the optimal set of provided waypoints or the optimal heading time-series. In both operational modes, the optimal path can always be followed in the environment flow realization for which it is truly optimal, i.e., for which it reaches the target \mathbf{x}_f in fastest time. For a vehicle programmed to achieve waypoints, the ability to traverse a non-optimal path depends on whether the currents are always weaker than the vehicle’s forward thrust or not. If this is the

case, then all prescribed paths are feasible in finite time and differ from the optimal path mainly in total travel time. If not, then the vehicle does not have local controllability and may get advected away from the desired path. Sometimes it may not be able to reach \mathbf{x}_f at all and thus needs to be aborted. If a vehicle programmed to achieve the optimal heading time-series for a given environment is operated in an environment for which these headings are sub-optimal, the vehicle will in general end up at a different end location than its intended target.

To quantify the physical dissimilarity between the paths $\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m)$ and $\mathbf{X}_P(\mathbf{x}_s, t; \omega_m)$, we employ the Fréchet distance δ_F [54,55] as the error metric Eq. (3a). A zero value of the Fréchet distance means that the paths are coincident and higher values progressively imply more dissimilarity. Alternative error metrics are possible for vehicles programmed with waypoint objectives or heading objectives. For example, for a vehicle programmed with waypoint objectives, the sub-optimal trajectory reaches the target in time $T(\mathbf{x}_f; \omega_l|\omega_m)$ instead of $T(\mathbf{x}_f; \omega_m)$. Thus, the error metric could be $|T(\mathbf{x}_f; \omega_l|\omega_m) - T(\mathbf{x}_f; \omega_m)|$ (Eq. (3b)). For a vehicle programmed with a heading objective, an appropriate error metric could be the error in arrival location. Specifically, following $\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m)$, the vehicle ends up at \mathbf{x}'_f instead of \mathbf{x}_f and the error metric is the distance $|\mathbf{x}_f - \mathbf{x}'_f|$ (Eq. (3c)). In summary, the above error metrics are

$$e_{lm} = \delta_F(\mathbf{X}_P(\mathbf{x}_s, t; \omega_l|\omega_m), \mathbf{X}_P(\mathbf{x}_s, t; \omega_m)), \tag{3a}$$

$$e_{lm} = |T(\mathbf{x}_f; \omega_l|\omega_m) - T(\mathbf{x}_f; \omega_m)|, \tag{3b}$$

$$e_{lm} = |\mathbf{x}_f - \mathbf{x}'_f|. \tag{3c}$$

Mathematically, all the above metrics defined in Eq. (3) have the property that if the path followed was the exact time-optimal path for the realized environment then the error is zero. For non-optimal paths, the error gets progressively higher with the degree of sub-optimality.

We note that the above metrics have been chosen based on our experience with real time missions with REMUS 600 AUVs [7,56,57]. Of course, other error metrics could be utilized, e.g., those based on the maximum background flow encountered along the path, or minimum bathymetry clearance along the path, or the number of times a vehicle has to perform a particular maneuver such as surfacing. Specific error metrics for individual missions and vehicles can be formulated based on an accurate performance characterization of the vehicle utilized, its navigation capabilities, fin configurations, and thresholds for satisfying navigation objectives. In the present paper, our goal is to provide a general theory and a recipe for risk-optimal path design and prediction for any acceptable metric e_{lm} . If a specific metric is well adapted to a given situation, one would simply need to update the definition of e_{lm} ; the rest of our schemes would still apply.

2.3.2. Analytical forms of the cost function f

The cost function f maps the error e_{lm} to a cost C_{lm} that specifies the utility cost associated with the choice l when the environment is m . The analytical form of the cost function f is important to characterize the risk profile of vehicle operators and appropriately compute the risk that will be utilized to order the choices [15]. For a risk-neutral preference of choice, the marginal increase in cost for a marginal increase in error should be relatively indifferent to the shape of the cost function. On the other hand, for a risk-seeking preference, the marginal increase in cost for a marginal increase in error should decrease monotonically with error, and for a risk-averse preference, the marginal increase in cost for a marginal increase in error should increase monotonically with error. Such a property ensures that for the former, higher errors are penalized less, whereas for the latter, higher errors are penalized more compared to a risk-neutral preference. In the rational utility theory, these risk-neutral, risk-seeking, and risk-averse preferences are modeled by utilizing linear, concave, and convex functions respectively [13,18,58].

In our applications we utilize the functions,

$$f(e_{lm}) = e_{lm}, \tag{4a}$$

$$f(e_{lm}) = \log_2(1 + e_{lm}), \tag{4b}$$

$$f(e_{lm}) = e_{lm} + 0.5 * e_{lm}^2, \tag{4c}$$

respectively to model risk-neutral, risk-seeking, and risk-averse preferences. Other linear, concave, and convex functions such as $f_{e_{lm}} = e_{lm}$, $e_{lm}^{0.5}$, and $-1 + 2^{e_{lm}}$ are also valid cost functions. To formally quantify the degree of risk aversion and choose an appropriate f , the Arrow–Pratt coefficient of absolute risk aversion applied on the errors, $A_r(e_{lm}) = f''(e_{lm})/f'(e_{lm})$ may be utilized [13], where the numerator measures the concavity and the denominator

Table 1

Risk-optimal path planning: Algorithm.

I. Probabilistic flow prediction

1. Predict the distribution of the velocity field $\mathbf{v}_{DO}(\mathbf{x}, t; \omega)$ by solving the discrete stochastic DO environmental flow equations (e.g., DO-NS Eq. (A.3) or DO-PEs [53]) to obtain the DO mean $\bar{\mathbf{v}}(\mathbf{x}, t)$, DO modes $\bar{\mathbf{v}}_i(\mathbf{x}, t)$, and the DO stochastic coefficients $\mu_i(t; r)$ for all $i = 1 \cdots n_{s,v}$ and $r = 1 \cdots n_r$.

II. Stochastic time-optimal path planning

1. Compute the distribution of the stochastic reachability fronts $\phi_{DO}(\mathbf{x}, t; \omega)$ by solving the stochastic DO level-set equations (Eq. (A.7)) with $\bar{\mathbf{v}}(\mathbf{x}, t)$, $\bar{\mathbf{v}}_i(\mathbf{x}, t)$, $\mu_i(t; r)$ for all $i = 1 \cdots n_{s,v}$ and $r = 1 \cdots n_r$.
2. Compute discrete time-optimal paths $\mathbf{X}_P(\mathbf{x}_s, t; r)$ and time-optimal headings $\hat{h}(t; r)$ for all $r = 1 \cdots n_r$ using the backtracking equation (Eq. (A.5)).

III. Risk evaluation and optimization

1. Simulate the trajectories $\mathbf{X}_P(\mathbf{x}_s, t; l|m)$ for waypoint objective $\mathbf{X}_P(\mathbf{x}_s, t; l)$ or heading objective $\hat{h}(t; l)$ for all $l = 1 \cdots n_r$.
2. Compute the error matrix e_{lm} as defined in Eq. (3) for all $l, m = 1 \cdots n_r$.
3. Compute the cost matrix $C_{lm} = f(e_{lm})$ corresponding to the cost functions as defined in Eq. (4) for all $l, m = 1 \cdots n_r$.
4. Compute the risk of choosing the path $\mathbf{X}_P(\mathbf{x}_s, t; l)$ as $R_l = \frac{1}{n_r} \sum_{m=1}^{n_r} C_{lm}$.
5. Find l that minimizes R_l . Then, return the risk optimal path $\mathbf{X}_P^*(\mathbf{x}_s, t) = \mathbf{X}_P(\mathbf{x}_s, t; l)$.

IV. Iterate

1. If there is evidence that the number of DO realizations $r = 1 \cdots n_r$ or adaptive sizes of the DO environmental or level-set subspaces, $n_{s,v}$ or $n_{s,\phi}$ Appendix A, are not sufficient to describe the probability density of $\mathbf{v}(\mathbf{x}, t; \bullet)$ or of $\mathbf{X}_P(\mathbf{x}_s, t; \bullet)$, augment the number of realizations or subspace sizes, and iterate.

normalizes this concavity. Such normalization is necessary as the utility representation is unique only up to affine transformations. For the same reason, the absolute value of the cost bears no significance to the risk optimization and only the relative cost of one choice with respect to another choice affects the risk optimization [59]. Therefore, in our applications (Section 3) we normalize the cost to lie between zero and one.

We note that an operator can assign specific cost values $f(e_{lm})$ and risk values (Eq. (1)) to specific realization mismatches e_{lm} . For example, the time-optimal vehicle for realization l could get caught in an unexpected current or wind field realization m of magnitude that is larger, at a certain place \mathbf{x} and time t , than the maximum vehicle speed. In such a situation, the optimal vehicle l is not controllable in the flow m . One may then simply utilize the above theory (Eq. (2)) since it applies to all situations and assign the cost as computed from the chosen function $f(e_{lm})$. This is what will be done here. Other more risk-averse operators could however assign specific higher cost and risk values to realization l . For example, they could replace $f(e_{lm})$ by the cost of the supremum of the possible finite error metric values e_{lm} . If they have no tolerance for such risk of unexpected lack of controllability (i.e. they abort the mission in such cases), they would directly eliminate these too risky optimal l from the samples considered in the risk-optimization (Eq. (2)).

2.4. Algorithm

The algorithm to implement the theory and schemes developed in Section 2.2–2.3 has four major steps as summarized in Table 1. See Table B.2 for our notation. A subscript *DO* is used to indicate that a DO decomposition has been applied to the corresponding stochastic variable.

In extreme cases, it is possible that after applying steps I to III to compute the risk-optimal path, further evidence emerges that the number of realizations n_r , and/or the subspace sizes $n_{s,v}$ or $n_{s,\phi}$ is not sufficient. For example, the probability distribution of the dynamic background flow might evolve such that more modes are required for improved accuracy or, the number of realizations needed to accurately describe the path probability distribution could be larger than the number of realizations required for the flow field. In such situations, one can increase the number of realizations or subspace sizes, and even iterate. Due to the nonlinear nature of the path planning equations, it is not always possible to simultaneously fix the optimal numbers of realizations or subspace sizes before computation. Adaptive schemes that update these numbers during the computation are then most useful [12,60,61].

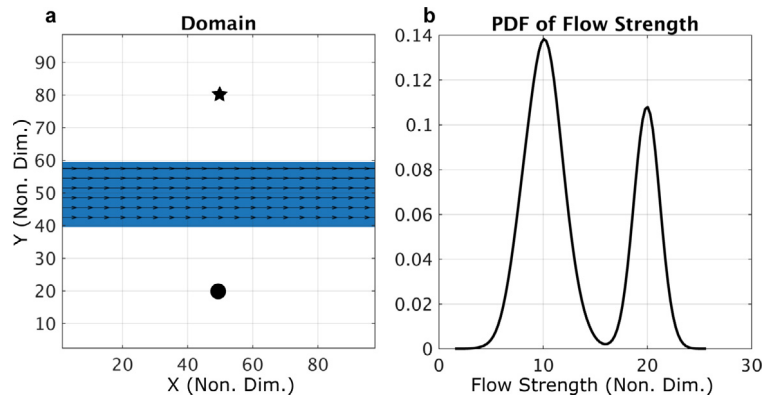


Fig. 2. Simulated stochastic steady-front crossing test case: Physical domain and non-Gaussian PDF of the flow strength: (a) In a square physical domain of non-dimensional side lengths 100×100 , an idealized stochastic steady front is modeled as a zonal-jet with uncertain strength, flowing from west to east between $y = 40$ and $y = 60$. (b) The PDF of the flow strength is a Gaussian Mixture Model with two Gaussians with non-dimensional mean, standard deviation, and mixture weight of $(10, 3, 0.65)$ and $(20, 1, 0.35)$, respectively.

3. Applications

In this section, we showcase the above theory and schemes by applying them to three environmental flows of increasing complexity. In the first, we consider an autonomous vehicle crossing a simulated stochastic steady front (in the ocean or atmosphere) similar to what we have used previously [1,2,4], but now with a bi-modal Gaussian Mixture Model (GMM) PDF for the uncertain flow strength and risk-optimal paths. In the second, we consider a stochastic quasi-geostrophic wind-driven double-gyre flow, as in [1], but now add the risk evaluation and optimization. In the third, we consider an autonomous vehicle navigating in an uncertain simulated flow exiting a strait/estuary (in the coastal ocean) or wind blowing through a widening constriction of an urban canopy (such as from a narrow street onto a wide street both lined with tall buildings). In all cases, the goal is to predict the risk-optimal path among the stochastic time-optimal paths. These test cases are pertinent not only for showcasing our methodology and for conducting numerical tests, but also for serving as canonical examples to analyze risk-optimal planning in realistic environmental flow scenarios in the future.

3.1. Stochastic steady front with uncertain flow strength

The first stochastic flow scenario we consider – a steady front with a non-Gaussian uncertain flow strength described by a GMM PDF – is an idealization of missions where an AUV or glider crosses a front or channel. Fig. 2a shows the domain and Fig. 2b shows the PDF of the flow strength. The domain is a non-dimensional square basin of size 100×100 with an uncertain jet flowing from east to west and confined between $40 < y < 60$. The direction of flow is considered known and only the strength is uncertain. The PDF of the flow strength is a GMM of two Gaussian components with non-dimensional mean, standard deviation, and mixture weights of $(10, 3, 0.65)$ and $(20, 1, 0.35)$ respectively. A single $n_{s,v} = 1$ mode is sufficient to describe the variability in the flow. We utilize $n_{r,v} = 5000$ realizations sampled from the GMM distribution of the stochastic coefficient corresponding to mode 1. Note that we consider a non-symmetric PDF for the flow strength (different from the uniform PDF used in [1]) in order to illustrate the key aspects of our risk-optimal planning. This completes step I of Table 1.

Next, we simulate the stochastic reachability and time-optimal path distribution for a vehicle with non-dimensional speed 20 starting at $(50, 20)$ and heading to a target at $(50, 80)$. The start point is depicted by a circular marker and the target by a star marker in all our figures. Eq. (A.7) is solved to obtain the distribution of stochastic reachability fronts. In Fig. 3a–e, five discrete snapshots during the evolution of the reachability front are shown. Each realization of the reachability front is colored by the jet strength of the corresponding flow realization. Fig. 3f shows the distribution of the time-optimal paths computed by solving Eq. (A.5). This completes step II of Table 1. Next, we compute the risk-optimal paths (step III of Table 1).

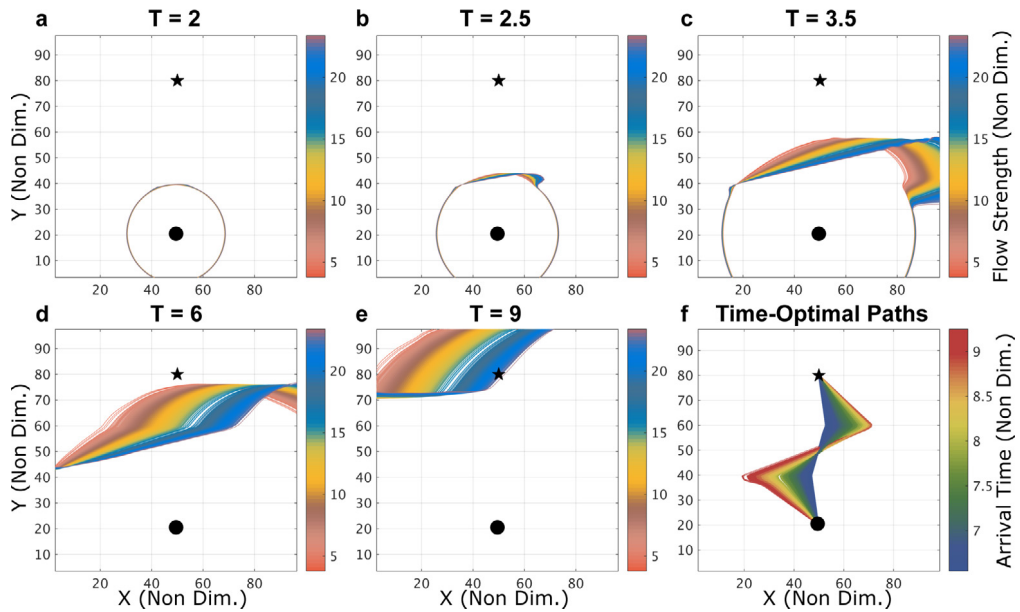


Fig. 3. Stochastic steady-front crossing: (a)–(e) Stochastic reachability front evolution at five discrete times colored by the non dimensional flow strength of that realization of the stochastic steady-front and (f) time-optimal path distribution for the vehicle crossing the steady front of uncertain flow strength, colored by the non dimensional arrival time at the target. In each panel, the circular black marker is the start location and the star black marker is the target.

3.1.1. Risk-optimal paths

Waypoint objective. For vehicles programmed to follow waypoint objectives, we simulate the trajectories $\mathbf{X}_P(\mathbf{x}_s, t; l|m)$ obtained by following the path $\mathbf{X}_P(\mathbf{x}_s, t; l)$ in the flow $\mathbf{v}(\mathbf{x}, t; m)$ for all $l, m = 1 \dots n_{r,\phi}$. This simulation is completed efficiently in parallel. The error metric matrix is then computed by Eq. (3) and normalized to lie between 0 and 1 (Section 2.3.2). The utility cost matrix C_{lm} for the required cost-functions can then be computed from e_{lm} by Eq. (4). This completes steps III.1 to III.3 of Table 1.

Fig. 4 row 1 shows the utility cost matrix for a risk-seeking cost function (column a), risk-neutral cost function (column b), and risk-averse cost function (column c). The l -axis of the cost matrix corresponds to the waypoint objective choice and the m -axis corresponds to the realizations of the flow field in which the choice l is utilized. The risk of the waypoint objective choices is then computed by marginalizing over m and obtaining the risk curve shown in row 2. The red point is the minimum risk choice. This completes steps III.4 and III.5 of Table 1). The realizations l are ordered according to the flow strength to aid visualization. For this bimodal GMM PDF of flow strength, the risk-seeking choice is the time-optimal path corresponding to the mean of the Gaussian with bigger mixture weight. The risk-neutral choice corresponds to a realization at the edge of the two mixtures and the risk-averse choice is a realization in the Gaussian with the smaller mixture weight. The risk curve also reflects the bi-modal nature of the flow strength PDF. In row 3, each waypoint objective choice is colored by its risk, and the risk-optimal trajectory (waypoint objective) is also shown in red, to match the red point in row 2.

The three risk-optimal waypoint objective choices corresponding to different risk behaviors are shown together in Fig. 5a. Physically, the risk-averse path is western most initially, followed by risk-neutral and risk-seeking. The more western the initial path is, the more the flow strength for that realization, see, e.g., [3,62]. The risk-seeking path is the true time-optimal path for a weaker flow strength than the risk-neutral path, which in turn is the true time-optimal path for a weaker flow strength than the risk-averse choice. The errors due to following these risk-optimal choices are shown by the PDF of e_{lm} for the optimal l in Fig. 5b. To recollect, the error e_{lm} is the Frechet distance between following the risk-optimal waypoint objectives (for the different cost functions) in all realizations of the stochastic flow environment and the true time-optimal path for that flow realization. Note that all the aborted missions have been assigned the maximum error value. The probability of the errors being low is highest for the risk-optimal path for a risk-seeking cost objective, but following this path also has the highest probability of the

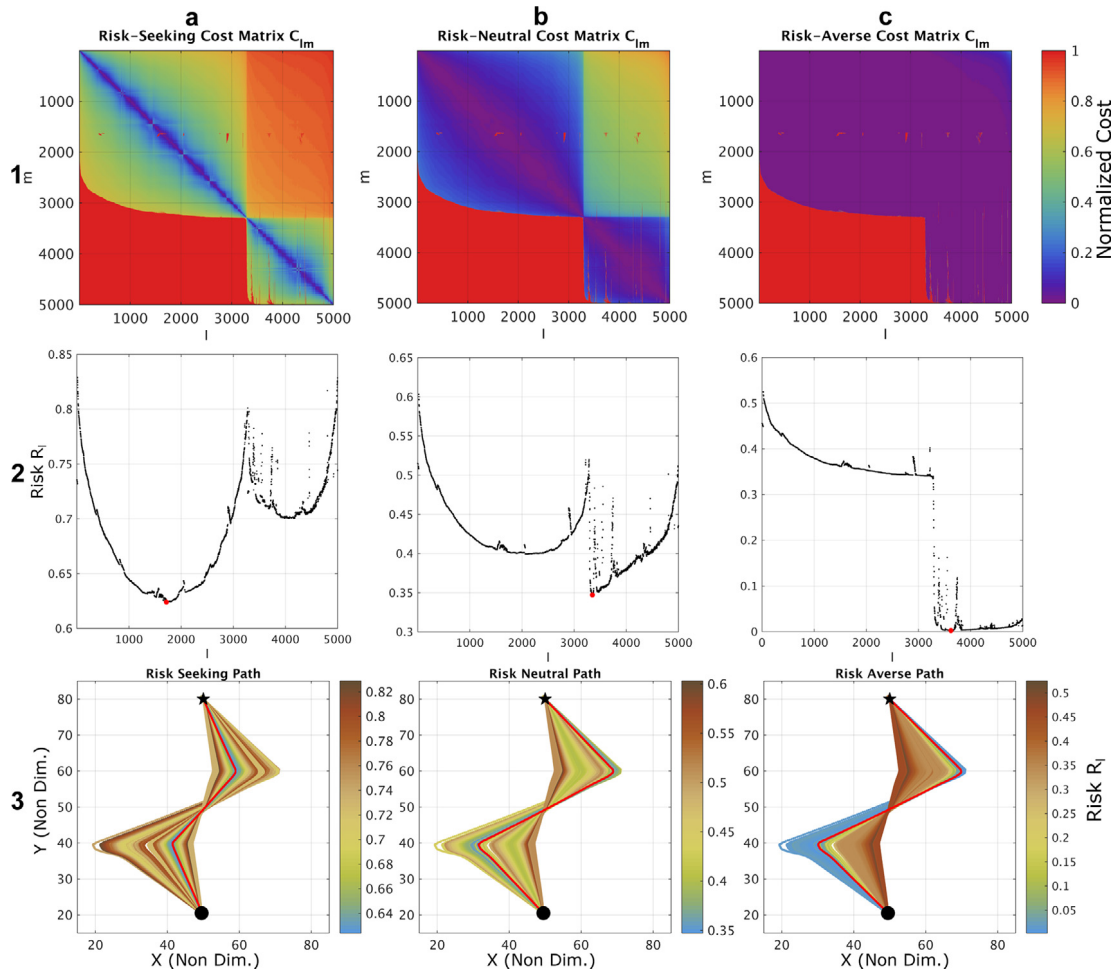


Fig. 4. Computation of risk-optimal paths with waypoint objectives for the stochastic steady-front crossing: Rows 1, 2, and 3 correspond to steps III.3, III.4, and III.5 respectively of Table 1. Columns a, b, and c correspond to risk-seeking, risk-neutral, and risk-averse behaviors, respectively. The cost matrix has been normalized (Section 2.3.1). To facilitate visualization, the waypoint objective choices l have been sorted, in rows 1 and 2, by the strength of the flow realization for which it is the exact time-optimal path. In row 3, each path is colored by the risk associated with following that waypoint objective. Here, the risk-optimal path is shown in a light shade of red; it corresponds to the dot in row 2 marked with the same shade of red.

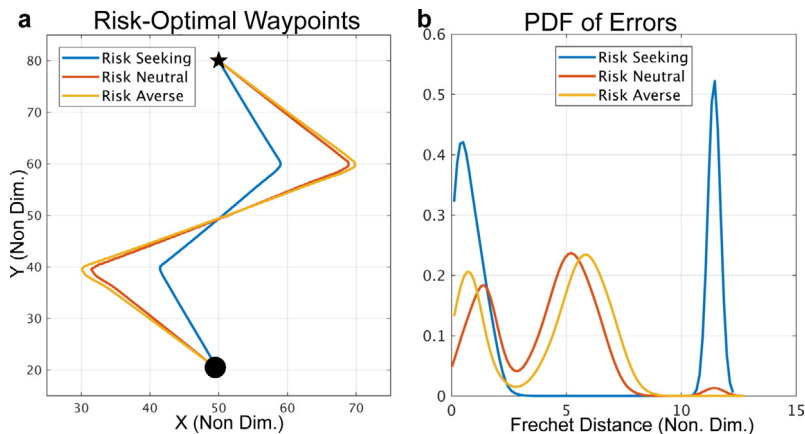


Fig. 5. Stochastic steady-front crossing: Risk-optimal paths and error distributions for the waypoint objectives. (a) The risk-seeking, risk-neutral, and risk-averse optimal paths for the waypoint objectives. (b) The PDF of error e_{lm} due to following the risk-optimal waypoint objective $\mathbf{X}_P(\mathbf{x}_s, t)$, quantifying the degree of sub-optimality due to following the risk-optimal paths.

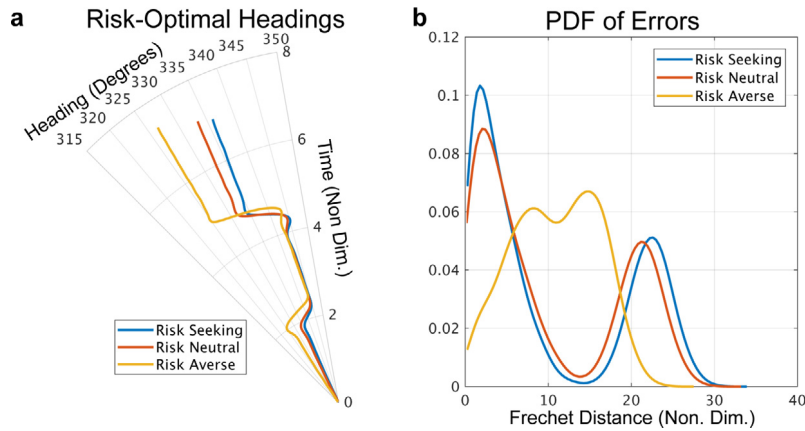


Fig. 6. Stochastic steady-front crossing: Risk-optimal headings and error distributions for the heading objectives. (a) The risk-seeking, risk-neutral, and risk-averse heading objective choices in a polar plot with time along the radial axis. (b) The PDF of errors e_{lm} due to following the risk-optimal heading objectives $\hat{h}(t; l)$.

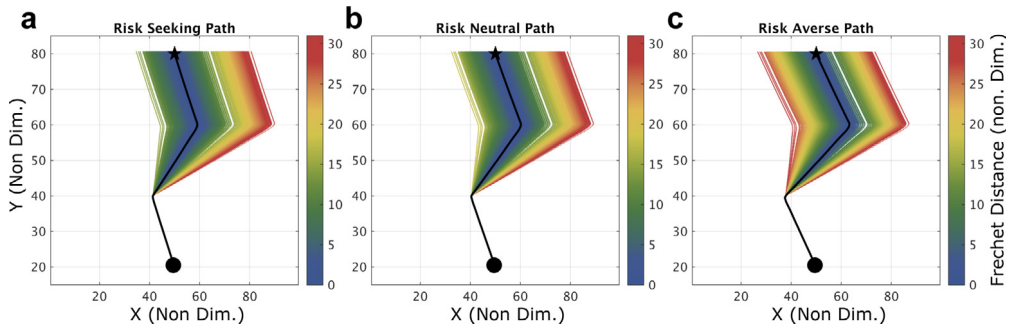


Fig. 7. Stochastic steady-front crossing: Error visualization (path distribution) for the risk-optimal heading objective choice. Each realized path corresponds to a particular flow realization and is colored by the discrete Frechet distance between that path and the true time-optimal path for the realized environmental flow. (a), (b), and (c) show the errors due to following risk-seeking, risk-neutral, and risk-averse choices, respectively. The PDF of these errors is shown in Fig. 6b.

mission being aborted. On the other hand, the probability of aborting the mission is the lowest for the risk-optimal path for a risk-averse cost function. The risk-neutral cost objective results in lower probability extreme errors and higher probabilities for errors with medium magnitude. Overall, depending on the risk tolerance of the operator, one of the risk-optimal choices will be appropriate. An aggressive operator may want to bet with the risk-seeking choice, whereas a conservative operator may choose the risk-averse choice and an ambivalent operator may choose the risk-neutral choice.

Heading objective. For vehicles programmed to follow heading objectives, we simulate the trajectories $\mathbf{X}_P(\mathbf{x}_s, t; l|m)$ obtained by following the heading objective $\hat{h}(t; l)$ in the flow $\mathbf{v}(\mathbf{x}, t; m)$ for all $l, m = 1 \dots n_{r,\phi}$. Similar to the waypoint objectives, steps III.2 to III.5 of Table 1 can be completed to obtain the risk-optimal heading objective choices. For brevity of presentation, we only show the risk-optimal heading choice and the error characteristics for these choices here.

Fig. 6a shows the computed risk-optimal heading choices for the three cost functions. The heading angles are represented on a polar plot with north being zero and angles increasing in the clockwise direction (east is 90, south 180, and west 270 degrees respectively). The radial axis is non-dimensional time. Fig. 6b shows the PDF of the error due to following the risk-optimal heading choice. The error is measured in the same way as for the waypoint objectives described above. The optimal heading for this jet-crossing test case has two changes, one on entering the jet and one on exiting, see, e.g., Lolla et al. [62] and Subramani et al. [3]. The initial heading is most western for the risk-averse path, followed by risk-neutral and then risk-seeking. The more western heading anticipates stronger




Fig. 8. Stochastic wind-driven double-gyre flow: The mean, standard deviation, and skewness of the velocity fields are shown in the first three rows, respectively. Perturbation from the DO mean of two representative realizations are shown in the last two rows. All fields are shown for three discrete times (columns), at the beginning, middle, and end of the planning horizon.

flow. In this scenario, as there is no flow except in $40 < y < 60$, the heading angle is exactly how the vehicle moves relative to the ground. Hence, we see that the optimal heading choices also mirror the characteristics of the

optimal waypoint choices. Following the risk-optimal heading choice in a flow for which it is sub-optimal results in the vehicle missing its target.

Fig. 7 shows the trajectories obtained by following the risk-optimal heading choices in all the flow realizations, colored by the Frechet distance of that path and the true time-optimal path for that flow. The risk-seeking heading objective choice results in paths that are close to the true time-optimal path in more flow realizations than the risk-neutral and risk-averse choices. On the flip side, the risk-seeking choice also leads to paths that are far from the true time-optimal paths in more flow realizations than the other two choices. As before, the risk-averse choice leads to more certain errors with medium magnitude.

Remark. If the flow strength for this test case was weaker than the relative vehicle speed and had a symmetric PDF, then all the risk-optimal paths coincide and correspond to the mean of the flow field (not shown here). In this case, the vehicle is fully controllable and all mission plans reach the target and no mission would get aborted. Additionally, in such cases, the trajectories $\mathbf{X}_P(\mathbf{x}_s, t; l|m) = \mathbf{X}_P(\mathbf{x}_s, t; l)$ and no new simulation is required for computing e_{lm} . Simulations and results with this assumption of full controllability are shown in [53].

3.2. Stochastic wind-driven double gyre

As our second illustration, we consider a stochastic wind-driven double gyre, ensuring continuity of our development [1]. The stochastic flow is obtained by solving the quasi-geostrophic DO equations (A.3), with $n_{s,v} = 5$ DO modes and $n_{r,v} = 5000$ DO realizations for 13.5 days. The domain is $1000 \text{ km} \times 1000 \text{ km}$ discretized into a 100×100 regular finite volume grid, with $\Delta x = \Delta y = 10 \text{ km}$. A deterministic zonal wind forcing is applied to drive the double gyre dynamics initialized with uncertain barotropic velocity components. Fig. 8 shows the mean, standard deviation, and skewness of the velocity fields at three discrete times $T = 0, 6.75, \text{ and } 13.5$ days. The perturbations from the DO mean of two extreme realizations are also shown, showing that significant perturbations in opposite directions are considered. These two extremes and the other 4998 realizations are, nonetheless, all simulated by one DO simulation using the DO mean, mode, and coefficient differential equations [1]. From the higher order statistics of the flow fields (e.g., the skewness on Fig. 8), we also see that the flow exhibits non-Gaussian behavior. This flow field and reachability fronts are described in [1]; here, the focus is on the risk-optimal planning.

3.2.1. Risk optimal paths

A mission with a vehicle moving at 40 cm/s is considered. As for the previous illustration, here we compute the risk-optimal paths for both waypoint objectives and heading objectives. We remark that the vehicle speed is smaller than the flow field. Hence, the vehicles are not locally fully controllable in this scenario and it is important to rigorously plan their paths, as done here utilizing exact S-PDEs.

Waypoint objectives. Results after completing step III for this test case is presented in Fig. 9, similar to Fig. 4. The rows 1 and 2 show the utility cost matrix and the risk curve with the minimum risk choice marked in red. The realizations l are ordered according to the DO coefficient 1 of the flow field to aid visualization. Row 3 shows all waypoint objective choices colored by their respective risk and the risk-optimal path colored in red, to match the red point in row 2. The columns a, b, and c correspond to the three risk behaviors: risk-seeking, risk-neutral and risk-averse, respectively.

Similar to the two panels of Fig. 5, Fig. 10a shows the three risk-optimal paths corresponding to the three different risk behaviors and Fig. 10b shows PDF of e_{lm} for the optimal l . Due to the strong flow in this test case, following the sub-optimal path has a high probability of aborting the mission. Here, the optimal paths for the different cost functions are physically close to each other. The optimal path for the risk-averse cost function has marginally better error characteristics than the other two. This illustration shows that our method can be utilized in complex stochastic flow scenarios and compute objective risk-optimal paths efficiently. It also shows that improvements in the forecast of the environmental condition, especially reducing the uncertainty in the velocity forecasts is essential for path planning in strong and dynamic flows. To compare the risk-optimal results when an assumption of complete controllability is made, we refer to Subramani [53].

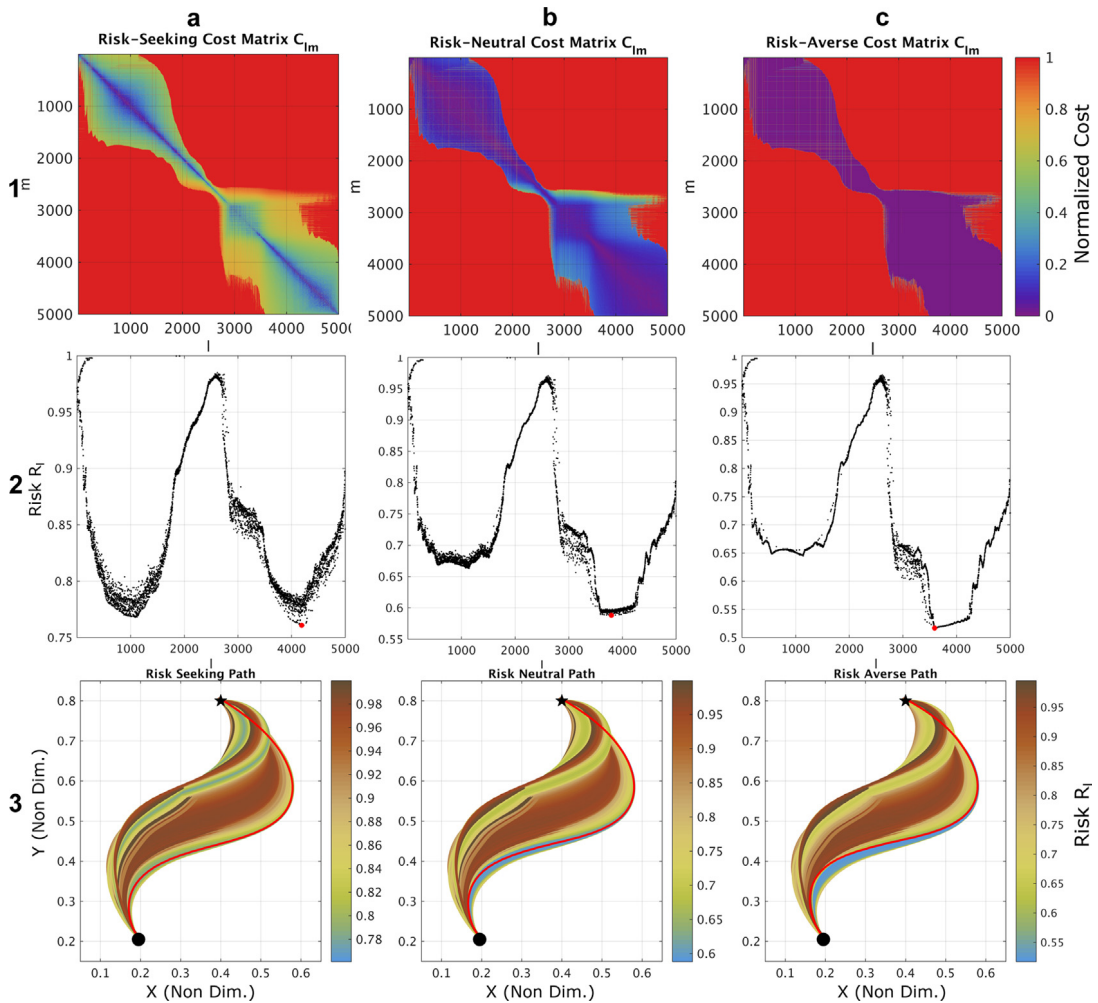


Fig. 9. Computation of risk-optimal paths with a waypoint objective for the stochastic wind-driven double-gyre flow: As Fig. 4, but for the mission in the stochastic double-gyre flow field. Here, to facilitate visualization, the waypoint objective choices l have been sorted, in rows 1 and 2, by the velocity DO Coefficient 1. In row 3, each path is colored by the risk associated with following that waypoint objective. Here, the risk-optimal path is shown in a light shade of red; it corresponds to the dot in row 2 marked with the same shade of red.

Heading objectives. Step III of Table 1 is completed for vehicles programmed to follow heading objectives in this stochastic double gyre flow field and final results are shown in Fig. 11, similarly to Fig. 6. Fig. 11a shows the risk-optimal heading objectives depicted as time-series of angles on a polar plot and Fig. 11b shows the PDF e_{lm} for the optimal l . The general direction of the heading is set by the mean flow and is along the north and north–north west for this test case. The flow variability is responsible for the deviations of the headings for different flow realizations. The optimal choices for the three cost functions are different with error characteristics and PDF shapes typical of each cost function. The bi-modal nature of the error PDF mirrors the PDF of the underlying flow field, see Subramani et al. [1].

In Fig. 12, we show the trajectories obtained by following the risk-optimal headings in all the flow realizations. Here, the paths are colored by the Frechet distance between the obtained trajectory and the true time-optimal path in that flow realization. For this flow field, the variability in the paths obtained by following the risk-optimal heading choice is wide around $y = 0.5$, but the end points for all trajectories lie in a small region around the target.

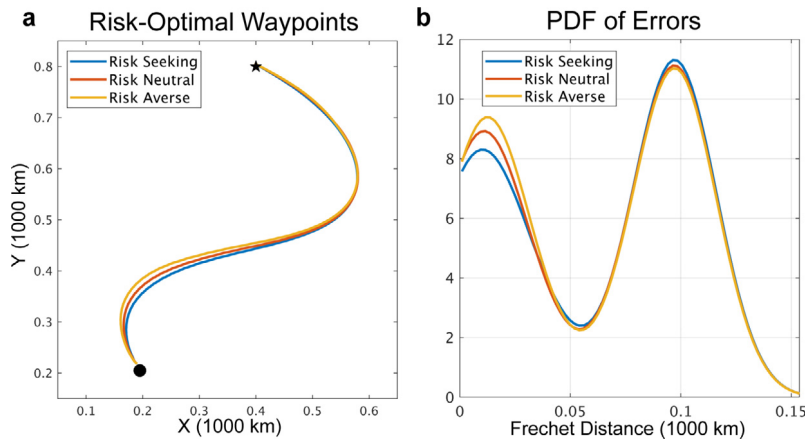


Fig. 10. Stochastic double-gyre flow: Risk-optimal paths and error distributions for the waypoint objectives. As Fig. 5 but for the mission in the stochastic double-gyre flow field.

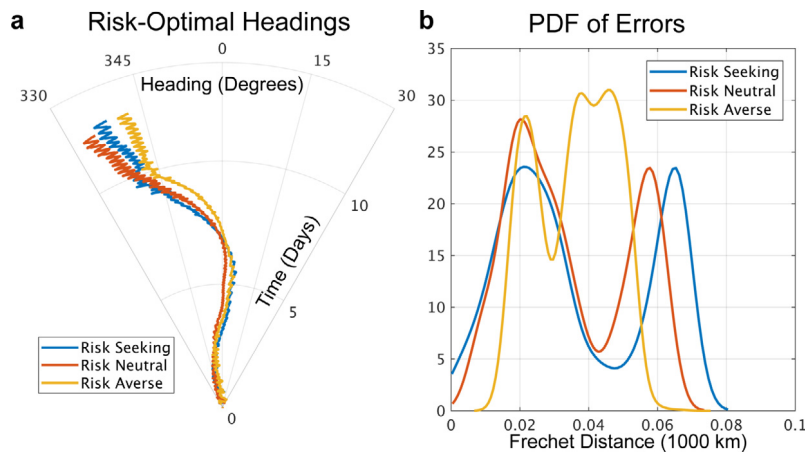


Fig. 11. Stochastic double-gyre flow: risk-optimal headings and error distributions for the heading objectives. As Fig. 6, but for the mission in the stochastic double-gyre flow field.

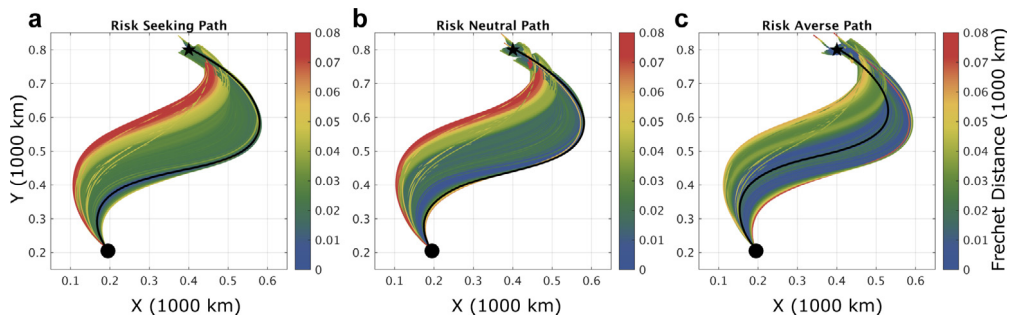


Fig. 12. Stochastic double-gyre flow: Error visualization (path distribution) for the risk-optimal heading objective choice. As in Fig. 7, but for the mission in the stochastic double-gyre flow field. Fig. 11b shows the PDF of these errors.

3.3. Stochastic flow exiting a strait

As our third illustration, we present risk-optimal planning in a stochastic flow scenario encountered in the coastal oceans and in urban environments. In the coastal oceans, a jet exiting a strait (or an estuary) into a wider channel creates eddies and meanders downstream [63]. In urban environments, the wind blowing through narrow constrictions between buildings into an open area also creates similar dynamics. Such flows can be idealized as sudden expansion flows studied extensively in fluid dynamics [64–66]. We consider a $6 \text{ km} \times 1 \text{ km}$ channel with a narrow inlet constriction of length $1/3 \text{ km}$ on the west. A uniform jet with a mean velocity of 50 cm/s is exiting the constriction into the wider channel. The initial conditions of the 2D-in-space barotropic velocity components in the channel are uncertain and sampled from a Gaussian covariance kernel with a decorrelation length scale of 1 km . The simulation is conducted by solving the quasi-geostrophic DO equations (A.3), with $n_{s,v} = 10$ DO modes and $n_{r,v} = 10,000$ DO realizations. The stochastic flow is allowed to develop for 1500 min by which time the flow forms recirculation zones and breaks to either the north or south of the centerline, depending on the initial uncertain perturbations. By that time, the nonlinear dynamics causes the initial Gaussian uncertainty to become non-Gaussian. For our risk-optimal planning, we consider the flow for the next 100 min .

Fig. 13 shows the mean field, the first three DO velocity modes, and the marginal PDF of the DO velocity coefficients at three discrete times $T = 0, 50, \text{ and } 100 \text{ min}$. As can be seen, the uncertainty in the flow field is highly dynamic and non-Gaussian, but the mean flow is nearly steady. Fig. 14 shows the standard deviation, skewness, and kurtosis of the velocity field at three discrete times. Also shown are two representative realizations with the jet breaking to the south or north. This completes step I of Table 1.

3.3.1. Stochastic reachability and time-optimal path distribution

Our path planning problem is to predict the risk optimal paths of a vehicle with a nominal relative speed of 25 cm/s traveling from the start point (depicted by a circle in Fig. 15) to three target locations A, B, and C (depicted by stars in Fig. 15). Using the velocity DO mean, modes, and coefficients computed above as inputs, we solve the stochastic DO level set equations (A.7) to compute the distribution of the stochastic reachability fronts. Fig. 15 shows the evolution of the stochastic reachability front at six discrete times, $T = 16, 33, 50, 66, 83 \text{ and } 100 \text{ min}$. In each panel, the reachability fronts are colored by the DO velocity coefficient 1 of the stochastic flow field. The strong inlet flow does not allow the vehicles to enter into the inlet constriction as can be seen by the reachability front not growing inside the inlet. The flow uncertainty creates the spread in the reachability front between $T = 33 \text{ min}$ and $T = 66 \text{ min}$, and hence there is significant variance in the arrival time at targets A and C. Thereafter, the spread in the reachability fronts reduces and between $T = 66 \text{ min}$ and $T = 100 \text{ min}$ the band of stochastic reachability fronts is narrow. The time-optimal path distribution is computed from the stochastic reachability fronts and backtracking equation. (A.5).

Fig. 16 shows the distribution of the time-optimal paths to the three targets, each colored with the arrival time at that target. In summary, the time-optimal paths to target B have high physical variability but low variability in the arrival time, whereas the paths to targets A and C have high variability in both the physical paths and arrival times. This completes step II of Table 1.

3.3.2. Risk optimal paths

For each of the three targets, we complete step III of Table 1 to obtain the risk optimal paths (for both waypoint and heading objectives, separately), similarly to the two previous test cases.

Waypoint objectives. As previously, we complete step III of Table 1 for waypoint objectives, for all three targets. Here, in Fig. 17, we show only the final risk-optimal paths and corresponding error PDFs for the three cost functions, and three targets. Column 1 shows the computed risk-optimal paths that minimize the risk-seeking, risk-neutral, and risk-averse cost functions for all the three targets (rows (a) to (c)). Column 2 shows the PDF of the errors due to following the computed risk-optimal waypoint objectives corresponding to the paths in column 1. Due to the complex and dynamic nature of the path uncertainty, the different risk-optimal paths show interesting behavior. For target A, the risk-seeking and risk-neutral paths are physically very close to each other and have similar error PDFs. They correspond to crossing or exiting the main uncertain jet, respectively (either errors can be small or large). For the risk-averse path, the errors are still bimodal, but of expected medium magnitude. For targets B and C, it is however the risk-neutral and risk-averse paths that are physically closer. This is because the risk-seeking

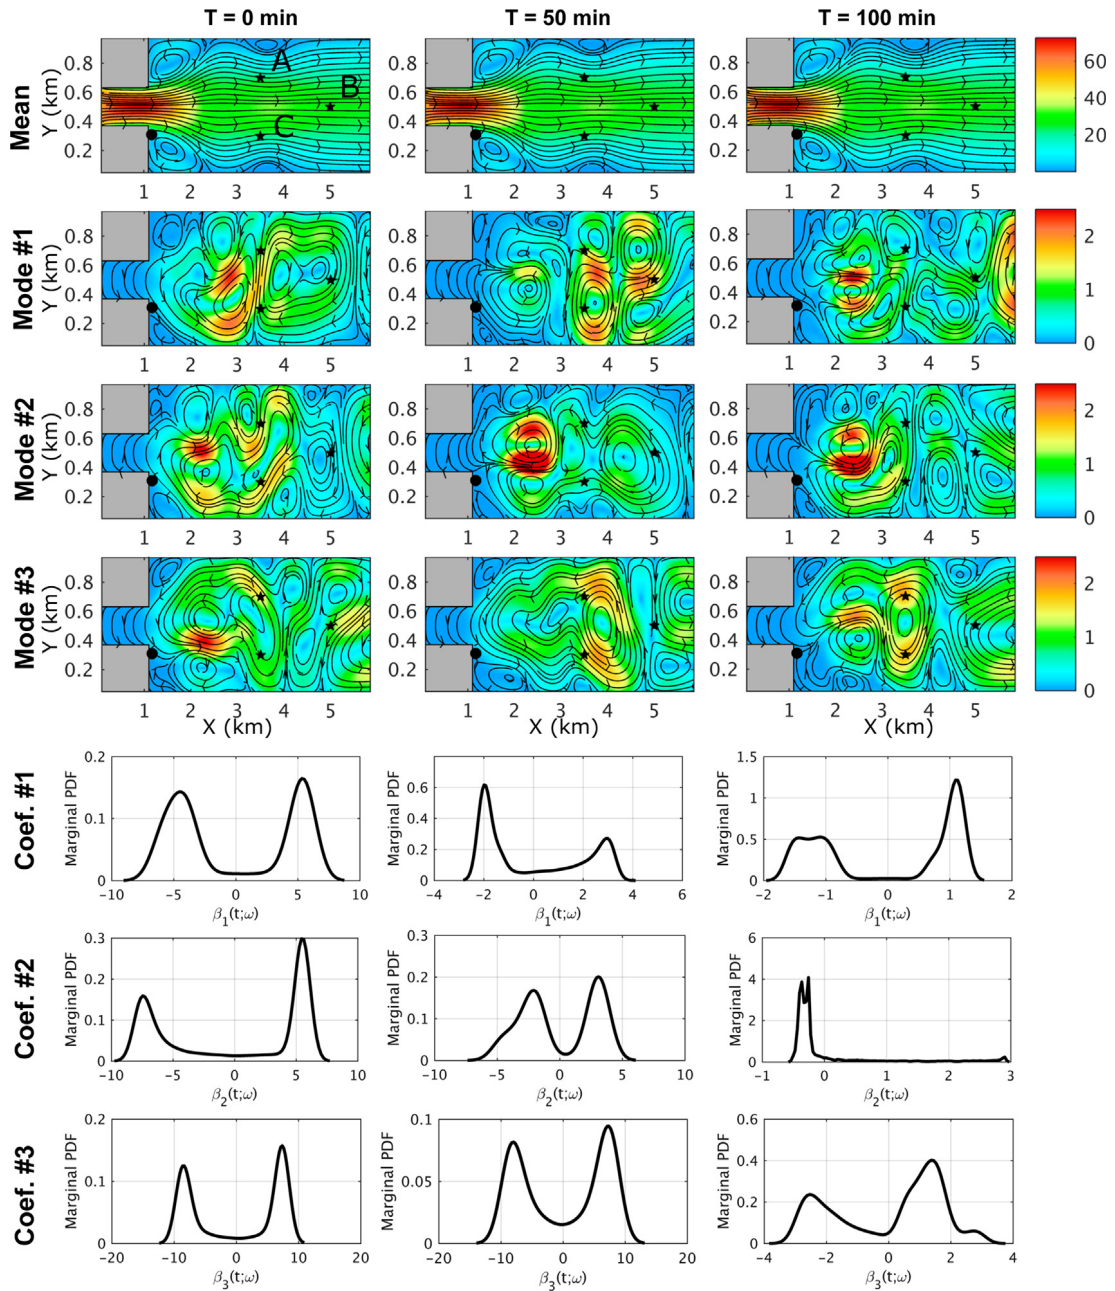


Fig. 13. Stochastic flow exiting a strait: Mean (in cm/s), Modes #1, #2, #3, and Coef. #1, #2, #3 are shown at three discrete times at the beginning, middle, and end of the planning horizon. (Top three rows): The mean and modes velocity streamlines are overlaid on a colorplot of their magnitude. (Bottom three rows): The coefficients are shown by their marginal PDF. The circular marker is the start point for the path planning mission and the star markers are the targets A, B, and C.

path “trusts” the uncertain jet and eddies much more, and so remains in the stronger current zones more. Its error distribution indeed shows the possibility of very large errors and/or mission abortion with relatively flat and wide PDF. For all targets, the shapes of the error PDFs of the three risk-optimal paths have similarities to the error PDF in other test cases. Crucially, our rigorous probabilistic prediction and risk-optimal path planning framework allows computation of such complex risk-optimal paths very efficiently.

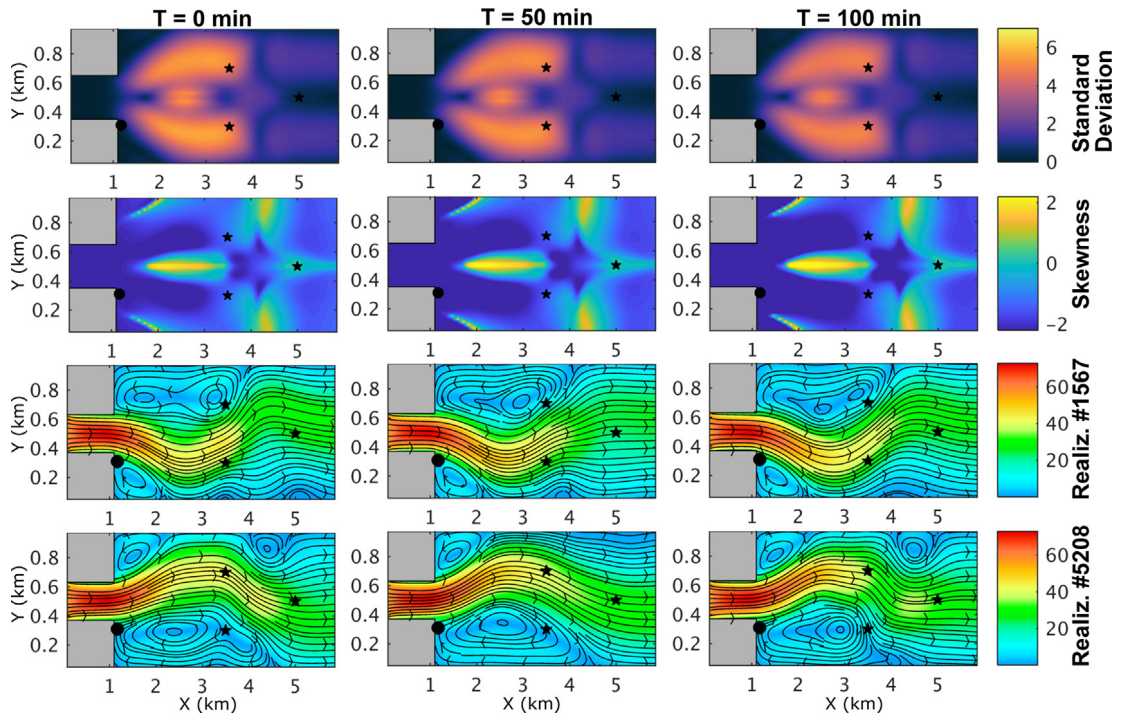


Fig. 14. Statistics and representative realizations of the stochastic flow exiting a strait: The standard deviation and skewness of the velocity fields are shown in the first two rows. Two representative realizations are shown in the last two rows. All fields are shown for three discrete times (columns), at the beginning, middle, and end of the planning horizon.

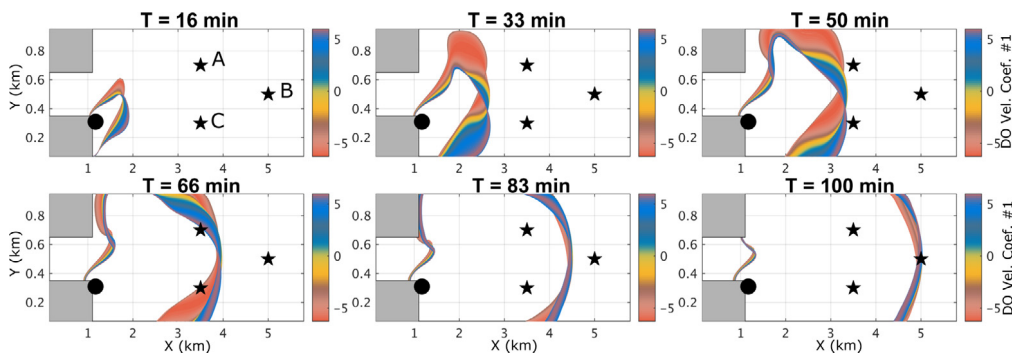


Fig. 15. Stochastic flow exiting a strait: Stochastic reachability from evolution colored by DO velocity coef. # 1. All reachability fronts are obtained by numerical integration of the stochastic DO level-set equations with the stochastic DO velocity fields of Fig. 14 as inputs. The circular marker is the start point for the path planning mission and the star markers are the targets A, B, and C.




Fig. 16. Stochastic flow exiting a strait: Time-optimal path distribution colored by the arrival time (in mins). The distribution of exact time-optimal paths computed from the stochastic reachability fronts using the stochastic backtracking equation.

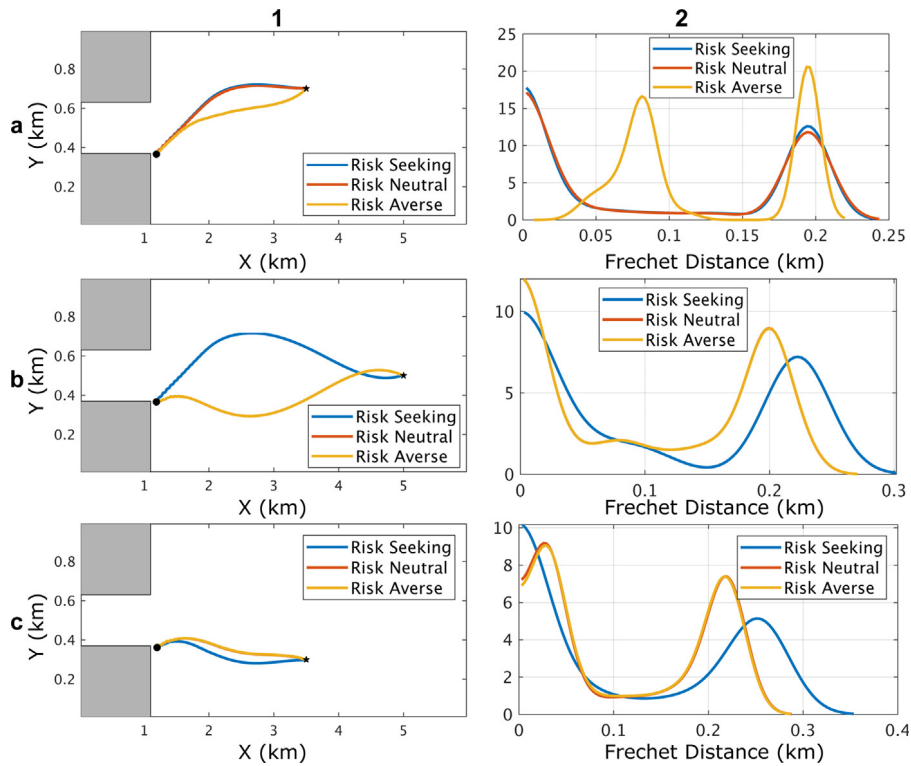


Fig. 17. Stochastic flow exiting a strait: Risk-optimal paths and error distributions for the waypoint objectives. As Fig. 5, but for all targets of the missions in the stochastic flow exiting a strait flow field. Rows (a)–(c) correspond to the three target locations. Column 1 shows the risk-seeking, risk-neutral, and risk-averse waypoint objective choices. Column 2 shows the PDF of errors e_{lm} due to following the risk-optimal waypoint objectives $X_P(x_s, t; l)$.

Heading objectives. The results after completing step III of Table 1 with heading objectives for this test case is shown in Fig. 18, similarly to Figs. 6 and 11 for the other two test cases. Targets A to C are shown in three columns, with the risk optimal headings and corresponding error PDFs in the two rows. For target A, the heading choices for the three cost functions form a time series with angles between north and east. Initially the three choices are similar, but the risk seeking choice changes the heading angle to near 60 degrees around $T = 30$ min. The risk-averse heading choice changes the angle to 60 degrees before 40 min and the risk neutral choice changes the angle after 40 min. After this abrupt change, all the three time series smoothly transition to 15 deg around $T = 60$ min. Target B lies in the mean strong flow, and the dynamic uncertainties make the headings vary significantly along the path to B. The risk-seeking choice anticipates that the flow is steadier than the other two risk profiles leading to high probability of extreme errors. The risk-averse choice forecasts the need to change the heading angles abruptly and its probability of extreme errors is low. The risk-neutral choice has a more gradual angle change along the time series and an error characteristic in between risk-seeking and risk-averse choices. For target C, the properties of the risk-optimal heading time series is similar to that of target A, but the general direction is to the south east. The error PDF shape and properties are analogous to those of the other cases.

Fig. 19 shows the trajectories obtained by following the risk-optimal heading objectives in all the flow realizations. As before, each trajectory is colored by the discrete Frechet distance between that and the true time-optimal path for that flow realization. The black path is obtained by following the heading objectives in the realization in which it is the true time-optimal heading time series. The three columns correspond to paths obtained by following the risk-averse, risk-neutral, and risk-seeking heading objectives to targets A, B, and C shown in the three rows. Here, as in the stochastic double gyre test case, the paths obtained by following the risk-optimal heading in all realizations have wide variability along the path, but the end points lie in a small region around the target.

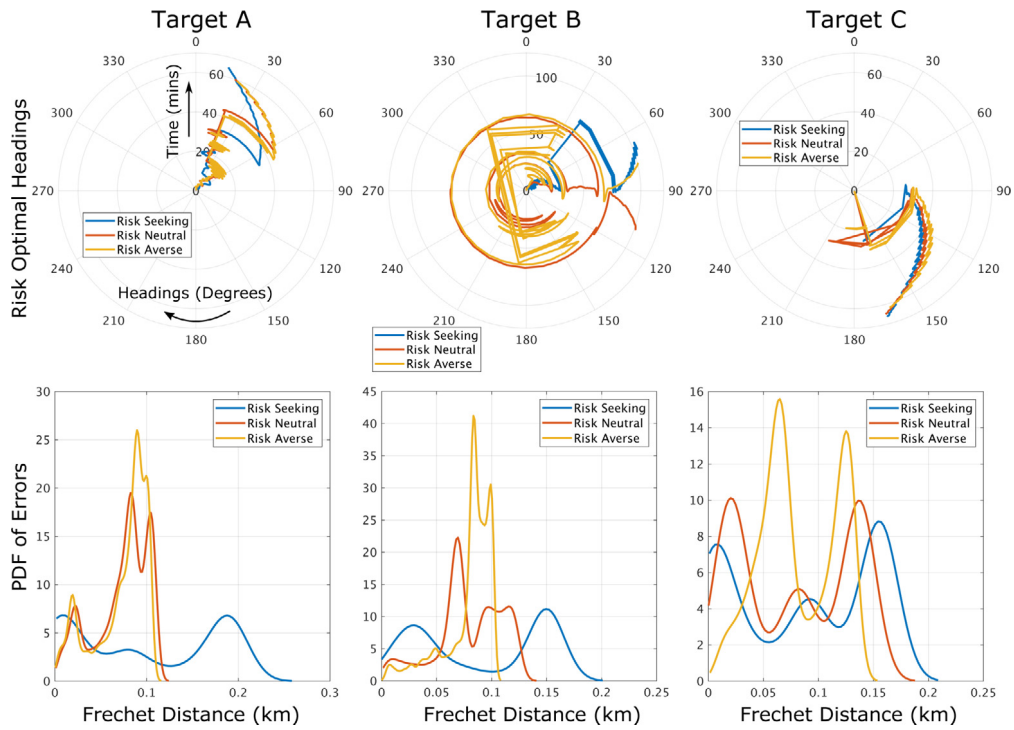


Fig. 18. Stochastic flow exiting a strait: Risk-optimal headings and error distributions for the heading objectives. As Fig. 6, but for all targets of the missions in the stochastic flow exiting a strait flow field. The top row shows the risk-optimal heading time-series choices in a polar plot with time along the radial axis. The bottom row shows the PDF of errors e_{lm} due to following the risk-optimal heading objectives $\hat{h}(t; l)$.

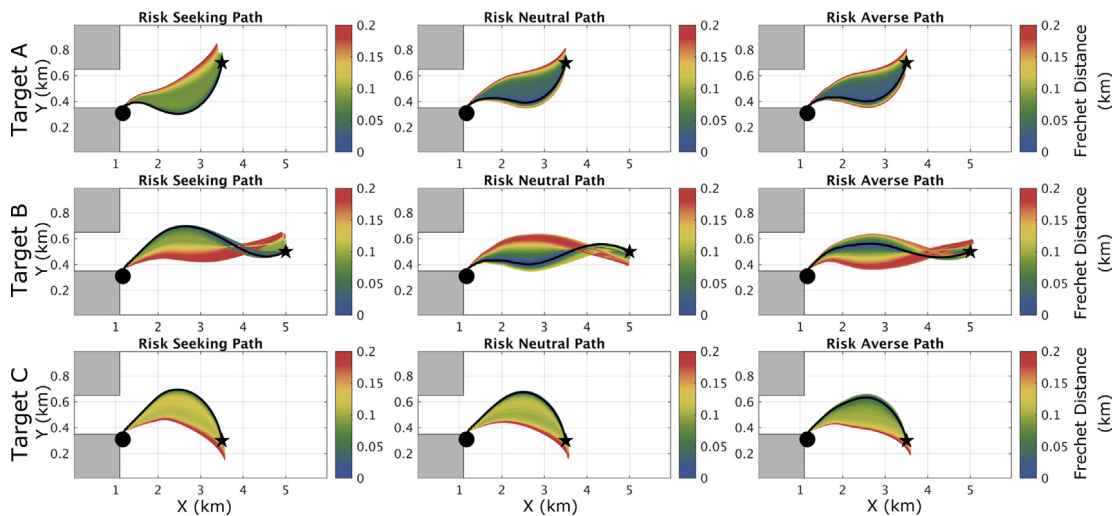


Fig. 19. Stochastic flow exiting a strait: Error visualization (path distribution) for the risk-optimal heading objective choice. As Fig. 7, but for all targets of the missions in the stochastic flow exiting a strait flow field. The rows represent the paths to the three targets A, B and C and the columns correspond to risk-seeking, risk-neutral, and risk-averse heading objective choices. Each trajectory is colored by the discrete Frechet distance between that trajectory and the exact time-optimal path corresponding to that realization of the environment.

4. Conclusion

We developed the theory for objective risk-optimal path planning by combining decision theory and our stochastic time-optimal path planning with stochastic DO level-set equations. The schemes and software developed compute risk-optimal paths for vehicles navigating in uncertain, strong and dynamic flows. Risk-optimal paths were defined as the time-optimal waypoint objectives or heading time-series objectives that minimize the expected utility cost of following a path that is possibly sub-optimal in the uncertain environment. Overall, the path planning proceeds in three steps: (i) obtain predictions of the probability distribution of environmental flows, (ii) obtain predictions of the distribution of exact time-optimal paths for the above flow distribution, and (iii) compute and minimize the risk of following the above time-optimal paths. To quantify and compute the risk, multiple error metrics were proposed including the discrete Frechet distance between following a prescribed path in a flow realization in which it is sub-optimal and the true time-optimal path for that flow realization. To convert the error into a utility cost, three functions corresponding to risk-seeking (concave), risk-neutral (linear) and risk-averse (convex) behaviors were utilized. The algorithm was presented and applied to plan risk-optimal paths in three illustrative stochastic flow scenarios: stochastic steady front crossing, double gyre, and flow exiting a strait.

In the first test case, the mission involved crossing a steady uncertain jet whose flow strength was a bimodal GMM PDF. The risk-seeking path was found to be the mean of the Gaussian component with higher mixture weight, the risk-neutral path was at the edge of the two Gaussians, and the risk-seeking path corresponded to a realization in the Gaussian component with lower mixture weight. In the second test case, the risk-averse path had the lowest probability of the mission getting aborted, followed by risk-neutral and risk-seeking paths. In the third test case, risk-optimal paths were planned for three targets. For one target, the risk-seeking and risk-neutral waypoint objectives were physically close whereas for the other two targets, it was the risk-averse and risk-neutral paths which were physically close to each other. In all scenarios, minimizing a risk-seeking cost function results in paths that have a higher probability of extreme errors, i.e., both low and high magnitudes. A risk-averse path on the other hand has a high probability of errors with medium magnitude, and a low probability of either extremes. Notably, in such complex flow situations, it is difficult to predict such behavior a priori; our methodology, however, allows the characterization and minimization of the risks along time-optimal paths for all flow conditions.

In the future, risk-optimal planning can be integrated with probabilistic current predictions [53,67] for real-time missions with real AUVs [7,43] or ships [23]. Machine learning techniques such as hierarchical clustering can be utilized for grouping massive ensembles of stochastic time-optimal paths, to find representative paths, and complete rapid risk-optimal calculations [68]. The theory, schemes, and software for rigorous risk-optimal computations may also have diverse applications in adaptive sampling [60,69], ocean acoustics [70–72], ecosystem management [53,73–75], and ocean energy [76], among others.

Acknowledgments

We thank C. Kulkarni and Q. Wei from our MSEAS group for useful discussions. We thank the MIT-Tata Center, USA for the fellowship support of DNS. We thank the Office of Naval Research, USA for research support under grants N00014-14-1-0476 (Science of Autonomy — LEARNS) and N00014-15-1-2616 (NASCar-OPS) to the Massachusetts Institute of Technology.

Appendix A. Summary of stochastic time-optimal path planning

A.1. Stochastic dynamically orthogonal barotropic quasi-geostrophic equations

In this Appendix, we briefly provide the stochastic PDEs that govern the dynamics of uncertain Barotropic Quasi-Geostrophic flows. For specifics, we refer to Ueckermann et al. [11]. The stochastic Barotropic Quasi-Geostrophic equation, written as conservation of momentum in the Langevin form is

$$\begin{aligned} \frac{\partial \mathbf{v}(\mathbf{x}, t; \omega)}{\partial t} + \nabla \cdot (\mathbf{v}(\mathbf{x}, t; \omega) \mathbf{v}(\mathbf{x}, t; \omega)) + f \hat{k} \times \mathbf{v}(\mathbf{x}, t; \omega) = \\ - \nabla p(\mathbf{x}, t; \omega) + \frac{1}{\text{Re}} \nabla^2 \mathbf{v}(\mathbf{x}, t; \omega) + a \boldsymbol{\tau}, \end{aligned} \quad (\text{A.1a})$$

$$\nabla \cdot \mathbf{v}(\mathbf{x}, t; \omega) = 0, \tag{A.1b}$$

$$\mathbf{v}(\mathbf{x}, 0; \omega) = \mathbf{v}_0(\mathbf{x}; \omega), \tag{A.1c}$$

where Re is the Reynold’s number, f is the Coriolis coefficient, and a the strength of the wind stress. For the Coriolis coefficient, we employ a β -plane approximation $f = f_0 + \beta y$.

Introducing a DO decomposition to $\mathbf{v}(\mathbf{x}, t; \omega)$,

$$\mathbf{v}(\mathbf{x}, t; \omega) = \bar{\mathbf{v}}(\mathbf{x}, t) + \sum_{j=1}^{n_{s,v}} \mu_j(t; \omega) \tilde{\mathbf{v}}_j(\mathbf{x}, t), \tag{A.2}$$

where $\mathbb{E}[\mu_j(t; \omega)] = 0, \forall j = 1 \dots n_{s,v}$, we can derive the DO mean (Eqs. (A.3a), (A.3b)), coefficient Eq. (A.3c), and mode equations (Eqs. (A.3d), (A.3e)) of the S-PDE Eq. (A.1) as

$$\nabla \cdot \bar{\mathbf{v}} = 0, \tag{A.3a}$$

$$\frac{\partial \bar{\mathbf{v}}}{\partial t} = \frac{1}{\text{Re}} \nabla^2 \bar{\mathbf{v}} - \nabla \cdot (\bar{\mathbf{v}} \bar{\mathbf{v}}) - C_{\mu_m \mu_n} \nabla \cdot (\tilde{\mathbf{v}}_m \tilde{\mathbf{v}}_n) - f \hat{k} \times \bar{\mathbf{v}} - \nabla \bar{p} + a \boldsymbol{\tau}, \tag{A.3b}$$

$$\begin{aligned} \frac{d\mu_i}{dt} = & \mu_m \left\langle \frac{1}{\text{Re}} \nabla^2 \tilde{\mathbf{v}}_m - \nabla \cdot (\tilde{\mathbf{v}}_m \bar{\mathbf{v}}) - \nabla \cdot (\bar{\mathbf{v}} \tilde{\mathbf{v}}_m) - \nabla \tilde{p}_m - f \hat{k} \times \tilde{\mathbf{v}}_m, \tilde{\mathbf{v}}_i \right\rangle \\ & - (\mu_m \mu_n - C_{\mu_m \mu_n}) \langle \nabla \cdot (\tilde{\mathbf{v}}_m \tilde{\mathbf{v}}_n), \tilde{\mathbf{v}}_i \rangle, \end{aligned} \tag{A.3c}$$

$$\nabla \cdot \tilde{\mathbf{v}}_i = 0, \tag{A.3d}$$

$$\frac{\partial \tilde{\mathbf{v}}_i}{\partial t} = \mathbf{D}_i^y - \langle \mathbf{D}_i^y, \tilde{\mathbf{v}}_k \rangle \tilde{\mathbf{v}}_k, \tag{A.3e}$$

where $\mathbf{D}_i^y = \frac{1}{\text{Re}} \nabla^2 \tilde{\mathbf{v}}_i - \nabla \cdot (\tilde{\mathbf{v}}_i \bar{\mathbf{v}}) - \nabla \cdot (\bar{\mathbf{v}} \tilde{\mathbf{v}}_i) - \nabla \tilde{p}_i - C_{\mu_i \mu_j}^{-1} M_{\mu_j \mu_m \mu_n} \nabla \cdot (\tilde{\mathbf{u}}_m \tilde{\mathbf{u}}_n)$,

where we have dropped the parenthesis $(\mathbf{x}, t; \omega)$ for brevity of notation. Here, \mathbb{E} is the expectation operator, C is the covariance matrix, M is the third moment, and $\langle \psi_i, \psi_j \rangle$ is the spatial inner-product operator for any two fields ψ_i and ψ_j defined as $\langle \psi_i, \psi_j \rangle \equiv \int_{\mathcal{D}} \psi_i \psi_j d\mathbf{x}$.

A.2. Stochastic dynamically orthogonal level-set equations

Here, we summarize the exact stochastic level-set equations and the efficient stochastic DO level-set equations derived for stochastic time-optimal path planning in [1]. The stochastic reachability front for a *vehicle-Q* (Fig. 1) is governed by the stochastic Hamilton–Jacobi (HJ) level-set equation

$$\frac{\partial \phi(\mathbf{x}, \mathbf{t}; \omega)}{\partial t} + F(t) |\nabla \phi(\mathbf{x}, t; \omega)| + \mathbf{v}(\mathbf{x}, t; \omega) \cdot \nabla \phi(\mathbf{x}, t; \omega) = 0, \tag{A.4}$$

with the initial condition $\phi(\mathbf{x}, 0; \omega) = |\mathbf{x} - \mathbf{x}_s|$ and, if needed, open boundary condition $\frac{\partial^2 \phi(\mathbf{x}, t; \omega)}{\partial \mathbf{n}^2} \Big|_{\delta \mathcal{D}} = 0$, where \mathbf{n} is the outward normal to the boundary $\delta \mathcal{D}$, ω is a random event, and ϕ is the reachability-front-tracking scalar level-set field (e.g., signed distance function). For every ω , the optimal arrival-time $T(\mathbf{x}_f; \omega)$ at \mathbf{x}_f is obtained by integrating Eq. (A.4) until the first time t such that $\phi(\mathbf{x}_f, t; \omega) \leq 0$. The optimal trajectory $\mathbf{X}_P(\mathbf{x}_s, t; \omega)$ is then given by the particle backtracking equation

$$\begin{aligned} \frac{d\mathbf{X}_P(\mathbf{x}_s, t; \omega)}{dt} = & -\mathbf{v}(\mathbf{X}_P(\mathbf{x}_s, t; \omega), t; \omega) - F(t) \frac{\nabla \phi(\mathbf{X}_P(\mathbf{x}_s, t; \omega), t; \omega)}{|\nabla \phi(\mathbf{X}_P(\mathbf{x}_s, t; \omega), t; \omega)|}, \\ 0 \leq t \leq & T(\mathbf{x}_f; \omega) \text{ and } \mathbf{X}_P(\mathbf{x}_s, T; \omega) = \mathbf{x}_f. \end{aligned} \tag{A.5}$$

To arrive at the DO level-set equations, let us introduce a DO decomposition to the stochastic level-set ϕ as

$$\phi(\mathbf{x}, t; \omega) = \bar{\phi}(\mathbf{x}, t) + \sum_{i=1}^{n_{s,\phi}} Y_i(t; \omega) \tilde{\phi}_i(\mathbf{x}, t), \tag{A.6}$$

where $\mathbb{E}[Y_i(t; \omega)] = 0 \forall i = 1 \dots n_{s,\phi}$.

Table B.2
Relevant notation.

Scalar		
i, j	$\in \mathbb{N}$	Stochastic subspace index
Q		Vehicle whose risk-optimal path is to be determined
F	$\in \mathbb{R}$	Vehicle speed
$n_{\mathcal{D}}$	$\in \mathbb{N}$	Dimension of \mathcal{D} , the physical domain
n_x	$\in \mathbb{N}$	Number of discrete grid points in the x -direction
n_y	$\in \mathbb{N}$	Number of discrete grid points in the y -direction
n_g	$\in \mathbb{N}$	Total number of discrete grid points in \mathcal{D}
$n_{s,\phi}$	$\in \mathbb{N}$	Dimension of the stochastic subspace of ϕ , the level-set scalar field
$n_{s,\mathbf{v}}$	$\in \mathbb{N}$	Dimension of the stochastic subspace of \mathbf{v} , the velocity vector field
Y_i	$\in \mathbb{R}$	Random variable describing the PDF of the orthonormal level-set (ϕ) modes $\tilde{\phi}_i$
μ_i	$\in \mathbb{R}$	Random variable describing the PDF of the orthonormal velocity (\mathbf{v}) modes, $\tilde{\mathbf{v}}_i$
n_r	$\in \mathbb{N}$	Number of Monte Carlo realizations of the level-sets (and velocity vector fields)
r	$\in \mathbb{N}$	Realization index
Vector		
\mathbf{v}	$\in \mathbb{R}^{n_{\mathcal{D}}n_g}$	Velocity vector field
$\bar{\mathbf{v}}$	$\in \mathbb{R}^{n_{\mathcal{D}}n_g}$	Mean velocity vector field
$\tilde{\mathbf{v}}_i$	$\in \mathbb{R}^{n_{\mathcal{D}}n_g}$	DO mode i of \mathbf{v} : Dynamically orthogonal basis for the stochastic subspace of \mathbf{v}
ϕ	$\in \mathbb{R}^{n_g}$	Level-Set field
$\bar{\phi}$	$\in \mathbb{R}^{n_g}$	Mean Level-Set field
$\tilde{\phi}_i$	$\in \mathbb{R}^{n_g}$	DO mode i of ϕ : Dynamically orthogonal basis for the stochastic subspace of ϕ

The stochastic DO level-set equations are

$$\frac{\partial \bar{\phi}}{\partial t} = - \left[F \mathbb{E}[\gamma] + \bar{\mathbf{v}} \cdot \nabla \bar{\phi} + C_{\mu_j Y_j} \tilde{\mathbf{v}}_j \cdot \nabla \tilde{\phi}_i \right], \tag{A.7a}$$

$$\frac{dY_i}{dt} = - \left\langle F(\gamma - \mathbb{E}[\gamma]) + Y_k \bar{\mathbf{v}} \cdot \nabla \tilde{\phi}_k + \mu_j \tilde{\mathbf{v}}_j \cdot \nabla \bar{\phi} + \tilde{\mathbf{v}}_j \cdot \nabla \tilde{\phi}_k (\mu_j Y_k - C_{\mu_j Y_k}), \tilde{\phi}_i \right\rangle, \tag{A.7b}$$

$$\frac{\partial \tilde{\phi}_i}{\partial t} = D_i - \left\langle D_i, \tilde{\phi}_n \right\rangle \tilde{\phi}_n, \tag{A.7c}$$

where $D_i = -C_{Y_j Y_i}^{-1} \left[F \mathbb{E}[Y_j \gamma] + C_{Y_j \mu_k} \tilde{\mathbf{v}}_k \cdot \nabla \bar{\phi} + \mathbb{E}[Y_j \mu_k Y_i] \tilde{\mathbf{v}}_k \cdot \nabla \tilde{\phi}_i \right] - \bar{\mathbf{v}} \cdot \nabla \tilde{\phi}_i$.

Appendix B. Notation

The notation for variables used in the present paper is provided here in Table B.2.

References

- [1] D.N. Subramani, Q.J. Wei, P.F.J. Lermusiaux, Stochastic time-optimal path-planning in uncertain, strong, and dynamic flows, *Comput. Methods Appl. Mech. Engrg.* 333 (2018) 218–237.
- [2] T. Lolla, P.F.J. Lermusiaux, M.P. Ueckermann, P.J. Haley Jr., Time-optimal path planning in dynamic flows using level set equations: Theory and schemes, *Ocean Dyn.* 64 (10) (2014) 1373–1397.
- [3] D.N. Subramani, T. Lolla, P.J. Haley Jr., P.F.J. Lermusiaux, A stochastic optimization method for energy-based path planning, in: S. Ravela, A. Sandu (Eds.), *DyDESS 2014*, in: LNCS, vol. 8964, Springer, 2015, pp. 1–12.
- [4] D.N. Subramani, P.F.J. Lermusiaux, Energy-optimal path planning by stochastic dynamically orthogonal level-set optimization, *Ocean Model.* 100 (2016) 57–77.
- [5] T. Lolla, P.J. Haley Jr., P.F.J. Lermusiaux, Time-optimal path planning in dynamic flows using level set equations: Realistic applications, *Ocean Dyn.* 64 (10) (2014a) 1399–1417.
- [6] D.N. Subramani, P.J. Haley Jr., P.F.J. Lermusiaux, Energy-optimal path planning in the coastal ocean, *J. Geophys. Res.: Oceans* 122 (2017) 3981–4003.
- [7] D.N. Subramani, P.F.J. Lermusiaux, P.J. Haley Jr., C. Mirabito, S. Jana, C.S. Kulkarni, A. Girard, D. Wickman, J. Edwards, J. Smith, Time-optimal path planning: Real-time sea exercises, in: *Oceans '17 MTS/IEEE Conference*, Aberdeen, 2017.

- [8] W. Sun, P. Tsiotras, T. Lolla, D.N. Subramani, P.F.J. Lermusiaux, Multiple-pursuer-one-evader pursuit evasion game in dynamic flow fields, *J. Guid. Control Dyn.* 40 (7) (2017).
- [9] W. Sun, P. Tsiotras, T. Lolla, D.N. Subramani, P.F.J. Lermusiaux, Pursuit-evasion games in dynamic flow fields via reachability set analysis, in: 2017 American Control Conference (ACC), Seattle, 2017, pp. 4595–4600.
- [10] T.P. Sapsis, P.F.J. Lermusiaux, Dynamically orthogonal field equations for continuous stochastic dynamical systems, *Physica D* 238 (23–24) (2009) 2347–2360.
- [11] M.P. Ueckermann, P.F.J. Lermusiaux, T.P. Sapsis, Numerical schemes for dynamically orthogonal equations of stochastic fluid and ocean flows, *J. Comput. Phys.* 233 (2013) 272–294.
- [12] F. Feppon, P.F.J. Lermusiaux, A geometric approach to dynamical model-order reduction, *SIAM J. Matrix Anal. Appl.* 39 (1) (2018) 510–538.
- [13] P.J. Schoemaker, The expected utility model: Its variants, purposes, evidence and limitations, *J. Econ. Literature* 52 (1982) 9–563.
- [14] L.G. Epstein, Behavior under risk: Recent developments in theory and applications, *Adv. Econ. Theory* 2 (1992) 1–63.
- [15] J. Von Neumann, O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton university press, Princeton, NJ, 2007.
- [16] K.J. Arrow, Utilities, attitudes, choices: A review note, *Econometrica* (1958) 1–23.
- [17] P.C. Fishburn, *Nonlinear Preference and Utility Theory*, Vol. 5, Johns Hopkins University Press, Baltimore, 1988.
- [18] M. LiCalzi, A. Sorato, The pearson system of utility functions, *European J. Oper. Res.* 172 (2) (2006) 560–573.
- [19] Y.K. Hwang, N. Ahuja, Gross motion planning – a survey, *ACM Comput. Surv.* 24 (3) (1992) 219–291.
- [20] S.M. LaValle, *Planning Algorithms*, Cambridge university press, Cambridge, UK, 2006.
- [21] J.-C. Latombe, *Robot Motion Planning*, Vol. 124, Springer Science & Business Media, New York, 2012.
- [22] A.A. Pereira, J. Binney, G.A. Hollinger, G.S. Sukhatme, Risk-aware path planning for autonomous underwater vehicles using predictive ocean models, *J. Field Robotics* 30 (5) (2013) 741–762.
- [23] G. Mannarini, N. Pardini, G. Coppini, P. Oddo, A. Iafrazi, Visir-i: small vessels–least-time nautical routes using wave forecasts, *Geosci. Model Dev.* 9 (4) (2016) 1597–1625.
- [24] B. Garau, A. Alvarez, G. Oliver, Path planning of autonomous underwater vehicles in current fields with complex spatial variability: an A* approach, in: *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, IEEE, 2005, pp. 194–198.
- [25] D. Rao, S.B. Williams, Large-scale path planning for underwater gliders in ocean currents, in: *Proceedings of Australasian Conference on Robotics and Automation*, 2009, pp. 28–35.
- [26] A. Chakrabarty, J. Langelaan, UAV flight path planning in time varying complex wind-fields, in: *American Control Conference*, IEEE, 2013, pp. 2568–2574.
- [27] M.P. Wellman, M. Ford, K. Larson, Path planning under time-dependent uncertainty, in: *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*. UAI'95, Morgan Kaufmann Publishers Inc, San Francisco, CA, USA, 1995, pp. 532–539, URL <http://dl.acm.org/citation.cfm?id=2074158.2074219>.
- [28] G. Kewlani, G. Ishigami, K. Iagnemma, Stochastic mobility-based path planning in uncertain environments, in: *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 1183–1189.
- [29] D. Kruger, R. Stolkin, A. Blum, J. Briganti, april. Optimal AUV path planning for extended missions in complex, fast-flowing estuarine environments, in: *Robotics and Automation, 2007 International Conference on*, IEEE, 2007, pp. 4265–4270.
- [30] J. Witt, M. Dunbabin, Go with the flow: Optimal auv path planning in coastal environments, in: J. Kim, R. Mahony (Eds.), *Proceedings of Australasian Conference on Robotics and Automation* (2008) 86–94.
- [31] D. Beylkin, Path Optimization for an Earth-Based Demonstration Balloon Flight, Caltech, Pasadena, CA, 2008.
- [32] A. Alvarez, A. Caiti, R. Onken, Evolutionary path planning for autonomous underwater vehicles in a variable ocean, *IEEE J. Ocean. Eng.* 29 (2) (2004) 418–429.
- [33] M.P. Aghababa, 3d path planning for underwater vehicles using five evolutionary optimization algorithms avoiding static and energetic obstacles, *Appl. Ocean Res.* 38 (2012) 48–62.
- [34] J.A. Sethian, Fast marching methods, *SIAM Rev.* 41 (2) (1999) 199–235.
- [35] C. Petres, Y. Pailhas, P. Patron, Y. Petillot, J. Evans, D. Lane, Path planning for autonomous underwater vehicles, *IEEE Trans. Robot.* 23 (2) (2007) 331–341.
- [36] M. Soullignac, P. Taillibert, M. Rueher, Time-minimal path planning in dynamic current fields, in: *IEEE International Conference on Robotics and Automation*, 2009, pp. 2473–2479.
- [37] D.R. Thompson, S.A. Chien, Y. Chao, P. Li, B. Cahill, J. Levin, O. Schofield, A.P. Balasuriya, S. Petillo, M. Arrott, M. Meisinger, Spatiotemporal path planning in strong, dynamic, uncertain currents, in: *Proceedings of IEEE International Conference on Robotics and Automation*, 2010, pp. 4778–4783.
- [38] J. Barraquand, J.-C. Latombe, A monte-carlo algorithm for path planning with many degrees of freedom, in: *Robotics and Automation, 1990. Proceedings., 1990 IEEE International Conference on*, IEEE, 1990, pp. 1712–1717.
- [39] T. Wang, O.P. Le Maître, I. Hoteit, O.M. Knio, Path planning in uncertain flow fields using ensemble method, *Ocean Dyn.* 66 (10) (2016) 1231–1251.
- [40] C. Amato, G. Konidaris, A. Anders, G. Cruz, J.P. How, L.P. Kaelbling, Policy search for multi-robot coordination under uncertainty, *Int. J. Robot. Res.* 35 (14) (2016) 1760–1778.
- [41] D. Kularatne, H. Hajieghrary, M.A. Hsieh, Optimal path planning in time-varying flows with forecasting uncertainties, in: *2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2018, pp. 1–8.
- [42] P.F.J. Lermusiaux, T. Lolla, P.J. Haley Jr., K. Yigit, M.P. Ueckermann, T. Sondergaard, W.G. Leslie, Science of autonomy: Time-optimal path planning and adaptive sampling for swarms of ocean vehicles, in: T. Curtin (Ed.), *Springer HandBook of Ocean Engineering: Autonomous Ocean Vehicles, Subsystems and Control*, Springer, 2016, pp. 481–498, ch21.

- [43] P.F.J. Lermusiaux, D.N. Subramani, J. Lin, C.S. Kulkarni, A. Gupta, A. Dutt, T. Lolla, P.J. Haley Jr., W.H. Ali, C. Mirabito, S. Jana, A future for intelligent autonomous ocean observing systems, *J. Mar. Res.* 75 (6) (2017) 765–813, the Sea. Volume 17, The Science of Ocean Prediction, Part 2.
- [44] J. Blythe, *An Overview of Planning under Uncertainty*, Springer Berlin Heidelberg, Berlin, Heidelberg, 1999, pp. 85–110.
- [45] A. Bry, N. Roy, Rapidly-exploring random belief trees for motion planning under uncertainty, in: *Robotics and Automation (ICRA)*, 2011 IEEE International Conference on, IEEE, 2011, pp. 723–730.
- [46] M. Boddy, T.L. Dean, *Solving Time-Dependent Planning Problems*, Brown University, Department of Computer Science, Providence, RI, 1989.
- [47] B. Burns, O. Brock, Sampling-based motion planning with sensing uncertainty, in: *Robotics and Automation*, 2007 IEEE International Conference on, IEEE, 2007, pp. 3313–3318.
- [48] J.K. Rosenblatt, Optimal selection of uncertain actions by maximizing expected utility, *Auton. Robots* 9 (1) (2000) 17–25.
- [49] L. De Filippis, G. Guglieri, F. Quagliotti, A minimum risk approach for path planning of uavs, *J. Intell. Robot. Syst.* 61 (1–4) (2011) 203–219.
- [50] E. Rudnick-Cohen, J.W. Herrmann, S. Azarm, Risk-based path planning optimization methods for unmanned aerial vehicles over inhabited areas, *J. Comput. Inf. Sci. Eng.* 16 (2) (2016) 021004.
- [51] M. Zabaranin, S. Uryasev, P. Pardalos, *Optimal Risk Path Algorithms*, US, Boston, MA, 2002, pp. 273–298, URL https://doi.org/10.1007/0-306-47536-7_13.
- [52] B. Miller, K. Stepanyan, A. Miller, M. Andreev, 3d path planning in a threat environment, in: *2011 50th IEEE Conference on Decision and Control and European Control Conference*, 2011, pp. 6864–6869.
- [53] D.N. Subramani, *Probabilistic Regional Ocean Predictions: Stochastic Fields and Optimal Planning* (Ph.D. thesis), Massachusetts Institute of Technology, Department of Mechanical Engineering, Cambridge, Massachusetts, 2018.
- [54] T. Eiter, H. Mannila, 1994. Computing discrete fréchet distance. Tech. rep. Technische Universitat Wien.
- [55] H. Alt, M. Godau, Computing the fréchet distance between two polygonal curves, *Int. J. Comput. Geometry Appl.* 5 (01n02) (1995) 75–91.
- [56] C. Mirabito, D.N. Subramani, T. Lolla, P.J. Haley Jr., A. Jain, P.F.J. Lermusiaux, C. Li, D.K.P. Yue, Y. Liu, F.S. Hover, N. Pulsone, J. Edwards, K.E. Railey, G. Shaw, Autonomy for surface ship interception, in: *Oceans '17 MTS/IEEE Conference*, Aberdeen, 2017.
- [57] J. Edwards, J. Smith, A. Girard, D. Wickman, D.N. Subramani, C.S. Kulkarni, P.J. Haley Jr., C. Mirabito, S. Jana, P.F.J. Lermusiaux, Data-driven learning and modeling of AUV operational characteristics for optimal path planning, in: *Oceans '17 MTS/IEEE Conference*, Aberdeen, 2017.
- [58] P.C. Fishburn, On the nature of expected utility, in: *Expected Utility Hypotheses and the Allais Paradox*, Springer, 1979, pp. 243–257.
- [59] M. Yildiz, *Microeconomic Theory III*. MIT OpenCourseWare. URL: <https://ocw.mit.edu/courses/economics/14-123-microeconomic-theory-iii-spring-2015/lecture-notes-and-slides/>, 2015.
- [60] P.F.J. Lermusiaux, Adaptive modeling adaptive data assimilation and adaptive sampling, *Physica D* 230 (1) (2007) 172–196.
- [61] T.P. Sapsis, P.F.J. Lermusiaux, Dynamical criteria for the evolution of the stochastic dimensionality in flows with uncertainty, *Physica D* 241 (1) (2012) 60–76.
- [62] T. Lolla, M.P. Ueckeremann, K. Yiğit, Haley, Jr, P. J, P.F.J. Lermusiaux, Path planning in time dependent flow fields using level set methods, in: *IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 166–173.
- [63] T. Lolla, P.J. Haley Jr., P.F.J. Lermusiaux, Path planning in multiscale ocean flows: coordination and dynamic obstacles, *Ocean Modell.* 94 (2015) 46–66.
- [64] W. Cherdron, F. Durst, J.H. Whitelaw, Asymmetric flows and instabilities in symmetric ducts with sudden expansions, *J. Fluid Mech.* 84 (1) (1978) 13–31.
- [65] F. Durst, A. Melling, J.H. Whitelaw, Low Reynolds number flow over a plane symmetric sudden expansion, *J. Fluid Mech.* 64 (1974) 111–128.
- [66] R.M. Fearn, T. Mullin, K.A. Cliffe, Nonlinear flow phenomena in a symmetric sudden expansion, *J. Fluid Mech.* 211 (1990) 595–608.
- [67] P.F.J. Lermusiaux, P.J. Haley Jr., S. Jana, A. Gupta, C.S. Kulkarni, C. Mirabito, W.H. Ali, D.N. Subramani, A. Dutt, J. Lin, A. Shcherbina, C. Lee, A. Gangopadhyay, Optimal planning and sampling predictions for autonomous and lagrangian platforms and sensors in the northern Arabian Sea, *Oceanography* 30 (2) (2017) 172–185, special issue on Autonomous and Lagrangian Platforms and Sensors (ALPS).
- [68] A. Dutt, D.N. Subramani, C.S. Kulkarni, P.F.J. Lermusiaux, Clustering of massive ensemble of vehicle trajectories in strong, dynamic and uncertain ocean flows, in: *OCEANS Conference 2018*, IEEE, Charleston, SC, 2018, pp. 1–7.
- [69] S.V.T. Lolla, *Path Planning and Adaptive Sampling in the Coastal Ocean* (Ph.D. thesis), Massachusetts Institute of Technology, Department of Mechanical Engineering, Cambridge, Massachusetts, 2016.
- [70] P.F.J. Lermusiaux, C.-S. Chiu, Four-dimensional data assimilation for coupled physical-acoustical fields, in: N.G. Pace, F.B. Jensen (Eds.), *Acoustic Variability*, 2002, Kluwer Academic Press, Saclantcen, 2002, pp. 417–424.
- [71] F.-P.A. Lam, P.J. Haley Jr., J. Janmaat, P.F.J. Lermusiaux, W.G. Leslie, M.W. Schouten, L.A. te Raa, M. Rixen, At-sea real-time coupled four-dimensional oceanographic and acoustic forecasts during Battlespace Preparation 2007, *J. Mar. Syst.* 78 (Supplement) (2009) S306–S320.
- [72] P.F.J. Lermusiaux, J. Xu, C.-F. Chen, S. Jan, L. Chiu, Y.-J. Yang, Coupled ocean–acoustic prediction of transmission loss in a continental shelfbreak region: Predictive skill, uncertainty quantification, and dynamical sensitivities, *IEEE J. Ocean. Eng.* 35 (4) (2010) 895–916.
- [73] Ş.T. Beşiktepe, P.F.J. Lermusiaux, A.R. Robinson, Coupled physical and biogeochemical data-driven simulations of Massachusetts Bay in late summer: real-time and post-cruise data assimilation, *J. Mar. Syst.* 40–41 (2003) 171–212.
- [74] G. Cossarini, P.F.J. Lermusiaux, C. Solidoro, Lagoon of Venice ecosystem: Seasonal dynamics and environmental guidance with uncertainty analyses and error subspace data assimilation, *J. Geophys. Res.: Oceans* 114 (C6) (2009).

- [75] P.F.J. Lermusiaux, P.J. Haley, W.G. Leslie, A. Agarwal, O. Logutov, L.J. Burton, Multiscale physical and biological dynamics in the Philippine Archipelago: Predictions and processes, *Oceanography* 24 (1) (2011) 70–89, Special Issue on the Philippine Straits Dynamics Experiment.
- [76] S. Deluca, S. Zanforlin, B. Rocchio, P.J. Haley, C. Foucart, C. Mirabito, P.F.J. Lermusiaux, Scalable coupled ocean and water turbine modeling for assessing ocean energy extraction, in: *OCEANS Conference 2018*, IEEE, Charleston, SC, 2018, pp. 1–8.