

ARTICLE TEMPLATE

An Extensible Framework for Probabilistic Search of Stochastically-moving Targets Characterized by Generalized Gaussian Distributions or Experimentally-defined Regions of Interest

Benjamin L. Hanson^a, Muhan Zhao^a, and Thomas. R. Bewley^a

^aDepartment of Mechanical and Aerospace Engineering, University of California San Diego, 9500 Gilman Dr, La Jolla, CA 92093

ARTICLE HISTORY

Compiled October 3, 2024

ABSTRACT

This paper presents a continuous-time framework for modeling the evolution of a probability density function (PDF) summarizing the region of interest (ROI) during the search for a stochastically-moving, statistically stationary target. This framework utilizes a Fokker-Planck partial differential equation representing the evolution of this PDF subject to: diffusion modeling the spread of the PDF due to the random motion of the target, advection modeling the relaxation of the PDF back to a specified steady profile summarizing the ROI in the absence of observations, and observations substantially reducing the PDF within the vicinity of the search vehicles patrolling the ROI. As a medium for testing the proposed search algorithm, this work defines a new, more general formulation for the multivariate Generalized Gaussian Distribution (GGD), an extension of the Gaussian Distribution described by shaping parameter β . Additionally, we define a formulation with enhanced flexibility, the Generalized Gaussian Distribution with Anisotropic Flatness (GGDAF). Two techniques are explored that convert a set of target location observations into a steady-state PDF summarizing the ROI of the target, wherein the steady-state advection is numerically solved for. This work thus provides a novel framework for the probabilistic search of stochastically-moving targets, accommodating both non-evasive and evasive behavior.

KEYWORDS

Probabilistic search; Fokker-Planck equation; generalized Gaussian distribution; anisotropic flatness

1. Introduction

Probabilistic search techniques (that is, methods that model uncertainty distributions when searching for targets, and route search vehicles accordingly) are essential for maximizing the effectiveness of search efforts. Approximately stated, such searches must appropriately balance the competing goals of maximizing the probability of finding the target quickly and ensuring that the target is found eventually. Techniques that focus on the discovery of stationary objects [1,2] update the uncertainty of a motionless target via observations by mobile agents, and consider cases where the observations

CONTACT Benjamin L. Hanson. Email: blhanson@ucsd.edu

CONTACT Muhan Zhao. Email: muz021@ucsd.edu

CONTACT Thomas R. Bewley. Email: tbewley@ucsd.edu

taken may be inaccurate (i.e., false positives and false negatives). Other techniques that consider stochastically-moving targets model the state of the target, rather than its probability distribution, and evolve that state randomly, then compare different searching strategies via a “probability of capture” metric that is approximated by a frequentist approach [3–5]. In contrast, the present work models the time evolution of the probability density function (PDF) representing the region of interest (ROI) via the Fokker-Planck partial differential equation (PDE), thus providing a representative model appropriate to guide the search for stochastically-moving targets. This PDF evolution procedure could be utilized as the base for other probabilistic search techniques where more detailed considerations are incorporated, such as detection inaccuracy, trajectory optimization, and optimal search termination conditions [6,7].

The framework we employ in this paper makes the assumption that the location of the target is constrained to a static domain space. To support the validity of this, we note that it is well documented that predators often hunt over defined territories, which they defend against the incursion of other predators [8]. Herbivores often forage over similarly defined territories, often referred to as their stable home range [9,10]. Furthermore, the natural behaviors of many organisms over specific temporal scales result in invariant probability measures, i.e. the stochastic location of the organism is a statistically stationary process [11,12]. In this paper, it is assumed that this phenomena, known as locational stationarity, accurately describes the ROI of the proposed target over the expected temporal scale. While the location of said target at any specific time may be unknown, and its movement over ROI modeled as “random”, information about this ROI may be used to effectively guide a search. Today, rarely is such information utilized to carry out a search optimally; consider, e.g., the \$100M F-35B aircraft that was lost by the US Marines in the woods of South Carolina on Sep 17, 2023 [13]. The consequential search for this lost, highly sensitive government asset, which lasted over 24 hours, involved multiple search vehicles flying overlapping “lawnmower” search trajectories (i.e., simple linear paths that go and and back over long stretches, as illustrated in Figure 1). This search pattern was far from optimal, potentially delaying the recovery of this highly sensitive asset by many hours, with potentially severe consequences.

This work considers the problem of searching for stochastically-moving, statistically stationary targets displaying either non-evasive behavior, where the presence of the search vehicles have no effect on the motion of the target, or evasive behavior, where the target is privy to the presence of the search vehicles, thus directly impacting its motion. In the search for the stochastically-moving target, observations that do not find the target along a certain search path suppress the likelihood of the target being near that search path for a while, but due to the motion of the target, the likelihood of the target being in that region relaxes back after a period of time after the search vehicles moves on; this relaxation of the ROI back towards $\bar{p}(\mathbf{x})$ is captured appropriately by the advection term in the Fokker-Planck PDE developed and tested in this work.

Our probabilistic search method can be summarized as follows: search vehicles are directed to traverse the ROI, taking (visual, infrared, audio, etc.) observations in the vicinity of their trajectories. At these vicinities, $\bar{p}(\mathbf{x})$ is reduced as a function of the acuity and scope of the search vehicle. As the purpose of our framework is demonstrating the evolution of the PDF, we assume that all observations fail to locate the target, and terminate the simulation after a set time, rather than waiting for discovery or a terminal condition. For non-evasive targets, the relaxation advection $\bar{\mathbf{v}}(\mathbf{x})$ is proportional to the rate at which the observations lose confidence; that is, the advection term forces the probability back to the “mowed down” spots where the search vehicles

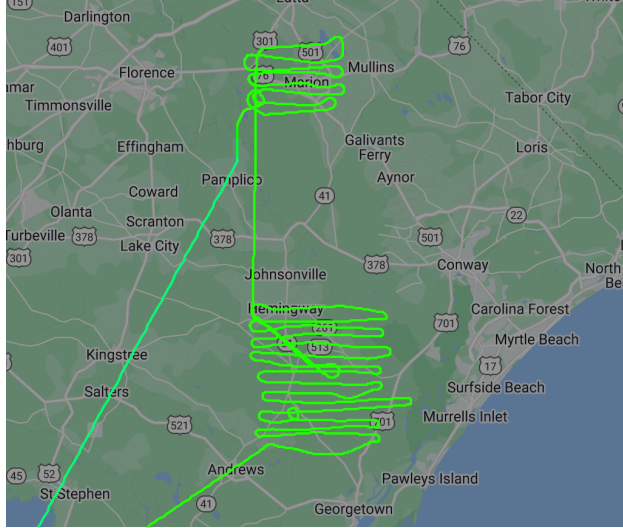


Figure 1. Search trajectory of CAP3935, one of the U.S. Air Force Civil Air Patrol planes tasked with searching for the lost F-35, from 2:17 PM UTC to 7:23 PM UTC on Sep 18, 2023 over South Carolina [14].

have recently been. We assume that the diffusion D is homogeneous and isotropic, i.e. the random fluctuations of the target are uncorrelated and constant with respect to space. For evasive targets, we model the evasive advection $\tilde{\mathbf{v}}(\mathbf{x}, t)$ as a function of the positions and disruptivities of the search vehicles, meant to represent the target being “scared off” by the encroaching vehicles, in addition to the relaxation advection $\bar{\mathbf{v}}(\mathbf{x})$. In this case, the diffusion of the target $D(\mathbf{x}, t)$ is no longer homogeneous, as the magnitude of “agitation” changes depending on the distances the search vehicles are from the target. Using a cooperative herding technique, we illustrate an effective method for converging the target probability density, thus increasing probability of discovery in the vicinity of the search vehicles.

To define the ROI of a target, we consider two scenarios: first, the ROI may be well-approximated by an analytical distribution, and second, numerous observations of said target may be aggregated to numerically define the ROI. In the analytical regime, we specifically consider the Gaussian Distribution and its lesser-known generalization. The Gaussian (a.k.a normal) Distribution (GD), originally defined by Gauss [15] in 1823, is often considered when representing the randomness of natural phenomena, characterized by a PDF of the form

$$p(\mathbf{x} | \boldsymbol{\mu}, \Sigma) = \frac{1}{(2\pi)^{\frac{d}{2}} |\Sigma|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\} \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$ is the random state vector, $\boldsymbol{\mu} \in \mathbb{R}^d$ is the mean of \mathbf{x} and $\Sigma \in \mathbb{R}^{d \times d}$ is the covariance of \mathbf{x} . It follows from this definition that the zeroth, first, and second central moments of the GD are

$$1 = \int_{\mathbb{R}^d} p(\mathbf{x} | \boldsymbol{\mu}, \Sigma) d\mathbf{x}, \quad \mathbb{E}[\mathbf{x}] = \boldsymbol{\mu}, \quad \text{and} \quad \mathbb{E}[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T] = \Sigma, \quad (2)$$

respectively.

The univariate generalized extension of the GD, with a shaping parameter β that controls the “peakedness” of the distribution, later dubbed the Generalized Gaussian Distribution (GGD), was established by Subbotin in 1923 [16]. The bivariate GGD PDF was introduced by Taguchi in 1978 [17], and later used for adaptive modeling in video coding in 1996 [18]. Both these formulations adhere to the central moment definitions given by (2). A multivariate extension of the GGD was then proposed by Gómez et al [19] in 1998 wherein the second central moment, or covariance, is no longer equal to the matrix Σ , thus diminishing the “generalization” denotation. Moreover, this specific multivariate GGD definition has been widely perpetuated by other sources in the literature and even utilized in the bivariate case over Taguchi’s definition [20–22]. In this paper, along with the primary objective of defining a novel probabilistic search technique, we aim to reformulate the multivariate GGD such that its covariance is the input matrix Σ , as is the case for the GD. We also introduce a novel distribution that allows for different shaping parameters along different eigenvectors of the covariance of the GGD, dubbed the Generalized Gaussian Distribution with Anisotropic Flatness (GGDAF).

The additional flexibility of the GGD is the shape parameter β , which may be utilized to more easily represent sharp peaks and flat tops. This idea is extremely pertinent in the fields of Gaussian Mixture Modeling (GMM) and Gaussian Sum Filtering (GSF) [23–26], wherein non-Gaussian uncertainty is approximated as a collection of Gaussian Distributions. Representing a nominal distribution with fewer (Generalized) Gaussians provides a distinct speedup in computation time. Generalized Gaussian Mixture Modeling (GGMM) and Generalized Gaussian Sum Filtering (GGSF) have been pursued [27–29], but rarely in multivariate form and never with our proposed formulation. By ensuring that the second central moment of the GGD is equal to the input matrix Σ , we theorize that the parameter estimation techniques that initialize GMM/GSF implementations, such as the Expectation-Maximization algorithm [30,31], may have closed-form update solutions for Σ , which is not currently the case. GGMM/GGSF is not pursued further in this paper.

2. Problem formulation

We focus in this paper on the class of problems in which the target of interest is stochastically-moving within a specified ROI $\Omega \in \mathbb{R}^d$ that is statistically stationary. In such problems, the distribution of likely locations of this target within this ROI, assuming no measurements are available, is modeled as a steady PDF $\bar{p}(\mathbf{x})$, i.e. the target displays locational stationarity. When performing a “probabilistic search”, search vehicles modify this PDF over time, suppressing the likelihood of discovery of the target near the path recently traveled by the search vehicles, assuming the target is not discovered by the search vehicles (otherwise, the search is over). This time-varying PDF is denoted $p(\mathbf{x}, t) \geq 0$ for all \mathbf{x} and t , with $\int p(\mathbf{x}, t) d\mathbf{x} = 1$ for all t , and $p(\mathbf{x}, t) \rightarrow 0$ for $|\mathbf{x}| \rightarrow \infty$ for any t .

In the case of non-evasively moving targets, the motion of the search vehicles does not change the motion of the target, and as time passes, away from the search vehicles, the PDF $p(\mathbf{x}, t) \geq 0$ relaxes back to its steady-state. This relaxation back to $\bar{p}(\mathbf{x})$ is achieved via the combined effects of advection field $\bar{\mathbf{v}}(\mathbf{x})$ and the symmetric diffusion

tensor $D(\mathbf{x}) \in \mathbb{R}^{d \times d} > 0$ in the Fokker-Planck PDE [32–34]

$$\frac{\partial p(\mathbf{x}, t)}{\partial t} + \sum_{i=1}^n \frac{\partial [v_i(\mathbf{x}) p(\mathbf{x}, t)]}{\partial x_i} = \sum_{i=1}^n \sum_{j=1}^n \frac{\partial^2 [D_{i,j}(\mathbf{x}) p(\mathbf{x}, t)]}{\partial x_i \partial x_j},$$

where $v_i(\mathbf{x})$ is the i^{th} component of the advection $\mathbf{v}(\mathbf{x})$ and $D_{i,j}(\mathbf{x})$ is the (i, j) th component of the diffusion tensor $D(\mathbf{x})$. This equation may be written in Gibbs notation as

$$\begin{aligned} \frac{\partial p(\mathbf{x}, t)}{\partial t} &= -\nabla_{\mathbf{x}} \cdot [\mathbf{v}(\mathbf{x}) p(\mathbf{x}, t)] + \nabla_{\mathbf{x}} \cdot [D(\mathbf{x}) \nabla_{\mathbf{x}} p(\mathbf{x}, t) + p(\mathbf{x}, t) \nabla_{\mathbf{x}} \cdot D(\mathbf{x})] \\ &= \nabla_{\mathbf{x}} \cdot [D(\mathbf{x}) \nabla_{\mathbf{x}} p(\mathbf{x}, t) + p(\mathbf{x}, t) \nabla_{\mathbf{x}} \cdot D(\mathbf{x}) - \mathbf{v}(\mathbf{x}) p(\mathbf{x}, t)], \end{aligned} \quad (3)$$

where $\nabla_{\mathbf{x}}$ is the vector differential operator with respect to \mathbf{x} .

We now set out to derive $\bar{\mathbf{v}}(\mathbf{x})$ as a function of $D(\mathbf{x})$ and $\bar{p}(\mathbf{x}, t)$ such that, applying the advection and diffusion of the Fokker-Planck PDE (3) to an arbitrary PDF relaxes it to $\bar{p}(\mathbf{x})$ (from hereon, we denote a variable relating to the process of relaxation with a bar, i.e. \bar{a}). Once derived, we provide example simulations of the relaxation advection derived for various steady-state distributions in action.

Steady-state is achieved when the RHS in (3) is zero [35]; it is thus seen that the advection field $\bar{\mathbf{v}}(\mathbf{x})$ that eventually relaxes the PDF $p(\mathbf{x}, t)$, governed by (3), towards a specified steady-state distribution $\bar{p}(\mathbf{x})$ is given by

$$\bar{\mathbf{v}}(\mathbf{x}) = \frac{1}{\bar{p}(\mathbf{x})} D(\mathbf{x}) \nabla_{\mathbf{x}} \bar{p}(\mathbf{x}) + \nabla_{\mathbf{x}} \cdot D(\mathbf{x}). \quad (4)$$

When D is homogeneous (not a function of \mathbf{x}), (4) simplifies to $\bar{\mathbf{v}}(\mathbf{x}) = \frac{1}{\bar{p}(\mathbf{x})} D \nabla_{\mathbf{x}} \bar{p}(\mathbf{x})$. For homogeneous isotropic diffusion $D = \lambda I$ where $\lambda \in \mathbb{R}$, a common special case, $\bar{\mathbf{v}}(\mathbf{x})$ can be defined as the gradient of a scalar potential function $\bar{\phi}(\mathbf{x})$, $\bar{\mathbf{v}}(\mathbf{x} | \lambda) = \nabla_{\mathbf{x}} \bar{\phi}(\mathbf{x} | \lambda)$. It follows that

$$\bar{\phi}(\mathbf{x} | \lambda) = \lambda \ln(\bar{p}(\mathbf{x})); \quad (5)$$

note that $\bar{\mathbf{v}}(\mathbf{x} | \lambda)$ is everywhere oriented orthogonal to the local isocontour of $\bar{\phi}(\mathbf{x} | \lambda)$, in the direction of steepest ascent.

3. The Gaussian Distribution (GD) and corresponding relaxation advection

The Gaussian Distribution, denoted $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \Sigma)$ [15], is defined by a PDF of the form

$$p(\mathbf{x} | \boldsymbol{\mu}, \Sigma) = \frac{1}{(2\pi)^{\frac{d}{2}} |\Sigma|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\} \quad (6)$$

where $\mathbf{x} \in \mathbb{R}^d$. The gradient of $p(\mathbf{x} | \boldsymbol{\mu}, \Sigma)$ is

$$\nabla_{\mathbf{x}} p(\mathbf{x} | \boldsymbol{\mu}, \Sigma) = -p(\mathbf{x} | \boldsymbol{\mu}, \Sigma) \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}), \quad (7)$$

and thus, by (4), the advection field $\bar{\mathbf{v}}(\mathbf{x} | \boldsymbol{\mu}, \Sigma)$ that drives the Fokker-Planck PDE (3) towards this PDF is

$$\bar{\mathbf{v}}(\mathbf{x} | \boldsymbol{\mu}, \Sigma) = -D \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}). \quad (8)$$

In the special case of $\Sigma = \sigma^2 I$, the GD and the corresponding relaxation advection becomes

$$p(\mathbf{x} | \boldsymbol{\mu}, \sigma) = \frac{1}{(2\pi)^{\frac{d}{2}} \sigma^d} \exp \left\{ -\frac{1}{2} \frac{(\mathbf{x} - \boldsymbol{\mu})^T (\mathbf{x} - \boldsymbol{\mu})}{\sigma^2} \right\} \quad (9a)$$

$$\text{and } \bar{\mathbf{v}}(\mathbf{x} | \boldsymbol{\mu}, \sigma) = -\frac{1}{\sigma^2} D (\mathbf{x} - \boldsymbol{\mu}). \quad (9b)$$




Figure 2. (left) $p(\mathbf{x} | \boldsymbol{\mu}, \Sigma)$ and (right) $\bar{\phi}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \lambda = 1)$ of a 2D Gaussian with $\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} 1 & 0.4 \\ 0.4 & 1 \end{bmatrix}$. For $p(\mathbf{x} | \boldsymbol{\mu}, \Sigma)$, the isocontour levels are spaced linearly from 0.01 to $\max\{p(\mathbf{x} | \boldsymbol{\mu}, \Sigma)\}$.

4. The Generalized Gaussian Distribution (GGD) and corresponding relaxation advection

We propose a new, more general multivariate Generalized Gaussian Distribution (GGD), denoted $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta)$, as defined by a PDF of the form

$$p(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta) = A_{\beta,d} \exp \left\{ - \left[B_{\beta,d} (\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right]^{\beta} \right\}, \quad (10a)$$

$$\text{where } A_{\beta,d} = \left(\frac{B_{\beta,d}}{\pi} \right)^{\frac{d}{2}} \cdot \frac{\Gamma(\frac{d}{2}) \beta}{\Gamma(\frac{d}{2\beta}) |\Sigma|^{\frac{1}{2}}} \quad \text{and} \quad B_{\beta,d} = \frac{\Gamma(\frac{d+2}{2\beta})}{d \Gamma(\frac{d}{2\beta})}, \quad (10b)$$

where again, $\mathbf{x} \in \mathbb{R}^d$. By defining $A_{\beta,d}$ and $B_{\beta,d}$ as such, the second central moment of the GGD is equal to the input matrix Σ , as is the case for the GD. The gradient of

$p(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta)$ is

$$\nabla_{\mathbf{x}} p(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta) = -2\beta p(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta) B_{\beta,d} \left[B_{\beta,d} (\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right]^{\beta-1} \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}), \quad (11)$$

and thus by (4), the advection field $\bar{\mathbf{v}}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta)$ that drives the Fokker-Planck PDE (3) towards this PDF is

$$\bar{\mathbf{v}}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta) = -2\beta D B_{\beta,d} \left[B_{\beta,d} (\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right]^{\beta-1} \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}). \quad (12)$$

Note that the $\beta = 1$ case reduces the GGD formulation in (10) to the Gaussian formulation in (6). As before, (10)-(12) simplify slightly in the special case of $\Sigma = \sigma^2 I$. Assume for the remainder of the GGD examples, shown in Figures 4-5, the statistics used are $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta)$, where $\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} 1 & 0.4 \\ 0.4 & 1 \end{bmatrix}$, and the diffusion tensor is $D = I$. For each $p(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta)$ plot, the isocontours are linearly spaced from 0.01 to $\max\{p(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta)\}$. As illustrated in Figures 2-4, the GGD, for $\beta > 1$, is flatter in




Figure 3. (left) $p(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta = 2)$ and (right) $\bar{\phi}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta = 2, \lambda = 1)$ of a 2D GGD.




Figure 4. (left) $p(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta = 3)$ and (right) $\bar{\phi}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta = 3, \lambda = 1)$ of a 2D GGD.

the immediate vicinity of $\mathbf{x} = \boldsymbol{\mu}$, and away from that point eventually falls off towards zero much faster than does the Gaussian Distribution. Note also from the colorbar




Figure 5. (left) $p(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta = 0.6)$ and (right) $\bar{\phi}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta = 0.6, \lambda = 1)$ of a 2D GGD.

in Figures 2-4 that the flatter top results in a smaller normalization coefficient, as is expected from (10). The smaller β , the more peaked the PDF is, and the larger β , the flatter the PDF is. In general, the GGD is a well-defined function with an additional tunable parameter β for more accurately representing ROI.

4.1. Higher-order central moments and smoothness of the GGD

To validate our multivariate GGD formulation, we look to the known higher-order central moments of the 1D case. The 1D k^{th} central moment of a scalar probability distribution $p(x)$ [36] is defined as

$$\mu_k = \mathbb{E}[(x - \mathbb{E}[x])^k] \triangleq \int_{\mathbb{R}} (x - \mathbb{E}[x])^k p(x) dx. \quad (13)$$

Having already defined the first and second central moments of the GGD ($\mu_1 = \mu$ and $\mu_2 = \sigma^2$), we look to clarify the third and fourth ones. The third central moment of the 1D GGD defined by (10) is

$$\mu_3 = \mathbb{E}[(x - \mu)^3] = \int_{-\infty}^{\infty} (x - \mu)^3 \left[A_{\beta,1} \exp \left\{ - \left[B_{\beta,1} \frac{(x - \mu)^2}{\sigma^2} \right]^\beta \right\} \right] dx \quad (14)$$

The third central moment is used to define the skewness,

$$\gamma_1 = \frac{\mu_3}{\sigma^3}. \quad (15)$$

Skewness is a measure of the asymmetry of a PDF, with positive skewness indication a more elongated right tail relative to the left. It is obvious that the integrand of (14) is odd, therefore, by construction, the skewness of the 1D GGD is

$$\gamma_1 = 0.$$

The fourth central moment of the 1D GGD is

$$\mu_4 = \mathbb{E}[(x - \mu)^4] = \int_{-\infty}^{\infty} (x - \mu)^4 \left[A_{\beta,1} \exp \left\{ - \left[B_{\beta,1} \frac{(x - \mu)^2}{\sigma^2} \right]^\beta \right\} \right] dx. \quad (16)$$

The fourth central moment is used to define the kurtosis,

$$\gamma_2 = \frac{\mu_4}{\sigma^4} - 3. \quad (17)$$

Kurtosis is a measure of the “peakedness” of a PDF, with negative kurtosis indicating less elongated tails than the Gaussian Distribution. We find that the kurtosis for the GGD as a function of β is

$$\gamma_2(\beta) = \frac{\Gamma(\frac{5}{2\beta})\Gamma(\frac{1}{2\beta})}{\left[\Gamma(\frac{3}{2\beta})\right]^2} - 3 \begin{cases} > 0, & \beta < 1 \\ = 0, & \beta = 1 \\ < 0, & \beta > 1 \end{cases}. \quad (18)$$

Note that the case when $\beta = 1$ returns the kurtosis of the Gaussian Distribution ($\gamma_2 = 0$), whereas $\beta = 0.5$ returns the kurtosis of the Laplacian distribution ($\gamma_2 = 3$). As $\beta \rightarrow \infty$, the GGD approaches a uniform distribution, so to validate (18), we numerically confirm that $\gamma_2(\beta)$ approaches the kurtosis of a uniform distribution ($\gamma_2 = -1.2$) as $\beta \rightarrow \infty$, shown in Figure 6.




Figure 6. $\gamma_2(\beta)$ where (left) $\beta = [0.1, 1]$ and (right) $\beta = [1, 50]$.

To further define the GGD, we analyze its smoothness. The GD is a C^∞ function, meaning it is infinitely differentiable, and all orders of derivatives are continuous. Using the 1D case of (11) we find the first derivative of the 1D GGD to be

$$\frac{\partial p(x|\mu, \sigma^2, \beta)}{\partial x} = -\frac{2\beta B_{\beta,1} p(x|\mu, \sigma^2, \beta)}{\sigma^2} \left[B_{\beta,1} \frac{(x-\mu)^2}{\sigma^2} \right]^{\beta-1} (x - \mu) \in \begin{cases} C^\infty, & \beta \geq 1 \\ C^0, & \beta < 1 \end{cases}. \quad (19)$$

It is obvious that the derivative is not smooth at $x = \mu$ when $\beta < 1$, so the GGD when $\beta < 1$ is of class C^0 , not C^∞ .

4.2. Generalized Gaussian Distributions with Anisotropic Flatness (GGDAF)

To enhance the flexibility of the GGD such that it may fit even more exotic data, we consider anisotropic flatness. In this context, what we mean by “anisotropic flatness” is that, along one eigenvector \mathbf{s}_i of the covariance matrix Σ , the distribution has a shaping parameter β_i and along another eigenvector \mathbf{s}_j , the distribution has a shaping parameter β_j , where β_i may not equal β_j if $i \neq j$. We define the 2D Generalized Gaussian Distribution with Anisotropic Flatness (GGDAF) as

$$\tilde{p}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \boldsymbol{\beta}) = \tilde{A}_{\boldsymbol{\beta},2} \exp \left\{ - \left[\sum_{i=1}^2 \left(\tilde{B}_{\beta_i,2} q_i^2 \right)^{\beta_i} + (B_{\beta_1,2})^{\beta_1} \delta_{\beta_1 \beta_2} \sum_{i=1}^{\beta_1-1} \binom{\beta_1}{i} (q_1^2)^{\beta_1-i} (q_2^2)^i \right] \right\} \quad (20a)$$

$$\text{where } \tilde{A}_{\boldsymbol{\beta},2} = \begin{cases} |\Sigma|^{\frac{1}{2}} \prod_{i=1}^2 A_{\beta_i,1} & \text{if } \beta_1 \neq \beta_2 \\ A_{\beta_1,2} & \text{otherwise} \end{cases}, \quad \tilde{B}_{\beta_i,2} = \begin{cases} B_{\beta_i,1} & \text{if } \beta_1 \neq \beta_2 \\ B_{\beta_i,2} & \text{otherwise} \end{cases}, \quad \mathbf{q} = LS^T(\mathbf{x} - \boldsymbol{\mu}), \quad (20b)$$

$$L = \sqrt{\Lambda^{-1}}, \quad \text{and } \Sigma = S\Lambda S^{-1}, \quad (20c)$$

where $\Sigma = S\Lambda S^{-1}$ is the eigen decomposition of the covariance matrix (as Σ is symmetric and positive semi-definite). The \tilde{A} and \tilde{B} coefficients are again chosen such that the definitions of the first and second central moments in (2) hold. To demonstrate that the GGDAF simplifies to the GGD when $\beta_1 = \beta_2$ and displays anisotropic flatness otherwise, we provide Figure 7. We note that the current formulation only returns circular isocontours for $\beta_1 = \beta_2$ when $\beta_1, \beta_2 \in \mathbb{N}^+$, but future work will attempt to address this such that the GGDAF is applicable and continuous for all real values of β .

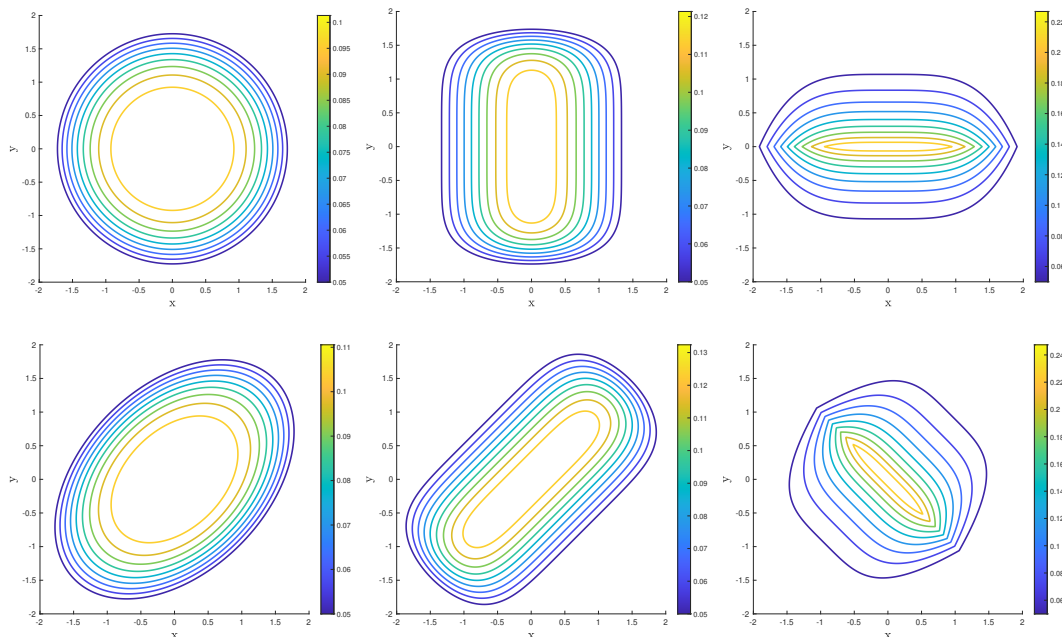


Figure 7. Examples of six, 2D GGDAF with parameters $\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$, (*top*) $\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ / (*bottom*) $\Sigma = \begin{bmatrix} 1 & 0.4 \\ 0.4 & 1 \end{bmatrix}$, and (*left*) $\boldsymbol{\beta} = [2 \ 2]$, (*middle*) $\boldsymbol{\beta} = [1 \ 3]$, and (*right*) $\boldsymbol{\beta} = [2 \ 0.5]$, respectively. In all subfigures, the isocontour levels are spaced linearly from 0.05 to $\max\{\tilde{p}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \boldsymbol{\beta})\}$.

For the 3D GGDAF we find the PDF to be

$$\tilde{p}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \boldsymbol{\beta}) = \tilde{A}_{\boldsymbol{\beta},3} \exp \left\{ - \left[\sum_{i=1}^3 \left(\tilde{B}_{\boldsymbol{\beta},3} q_i^2 \right)^{\beta_i} + \sum_{i=1}^3 \sum_{j>i}^3 \left\{ \left(\tilde{B}_{\boldsymbol{\beta},3} \right)^{\beta_i} \delta_{\beta_i, \beta_j} \sum_{k=1}^{\beta_i-1} \binom{\beta_i}{k} (q_i^2)^{\beta_i-k} (q_j^2)^k \right\} \right. \right. \\ \left. \left. + (B_{\boldsymbol{\beta},3})^{\beta_1} \delta_{\beta_1, \beta_2 \beta_3} \sum_{\substack{i,j,k \neq 0 \\ i+j+k=\beta_1}} \binom{\beta_1}{i,j,k} (q_1^2)^i (q_2^2)^j (q_3^2)^k \right] \right\} \quad (21a)$$

$$\text{where } \tilde{A}_{\boldsymbol{\beta},3} = \begin{cases} |\Sigma| \prod_{i=1}^3 A_{\beta_i,1} & \text{if } \beta_1 \neq \beta_2 \neq \beta_3 \\ |\Sigma|^{\frac{1}{2}} A_{\beta_i,1} A_{\beta_j,2} & \text{if } \beta_i \neq \beta_j = \beta_k \\ A_{\beta_i,3} & \text{otherwise} \end{cases}, \quad \tilde{B}_{\boldsymbol{\beta},3} = \begin{cases} B_{\beta_i,1} & \text{if } \beta_i \neq \beta_j, \beta_k \\ B_{\beta_i,2} & \text{if } \beta_i = \beta_j \neq \beta_k \\ B_{\beta_i,3} & \text{otherwise} \end{cases}, \quad \mathbf{q} = LS^T(\mathbf{x} - \boldsymbol{\mu}), \quad (21b)$$

$$L = \sqrt{\Lambda^{-1}}, \text{ and } \Sigma = S \Lambda S^{-1}. \quad (21c)$$




Figure 8. A 3D GGDAF with parameters $\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$, $\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, and $\boldsymbol{\beta} = [2 \ 1 \ 1]$. The isosurfaces on the (left) yz -plane are circular, as to be expected because $\beta_y = \beta_z$, while the (middle) xy - and (right) xz -planes demonstrate anisotropic flatness created by the fact that β_x is different than β_y and β_z . The three isosurface values shown are $p = 0.01$, $p = 0.001$, and $p = 0.0001$.

We provide an example of a 3D GGDAF in Figure 8 where $\beta_x \neq \beta_y = \beta_z$ to demonstrate circular isocontours in the xy -plane and noncircular isocontours in the other two perpendicular planes. The d -dimensional case of the GGDAF follows by splitting the d -nomial expansion into the subsequent d -nomial, $(d-1)$ -nomial, ..., trinomial, and binomial expansions, as demonstrated in the transition from (20) to (21), *mutatis mutandis*. Additionally, we note that, via Bayes' theorem [37],

$$p_{\mathbf{x}}(\mathbf{x}', t_{k+}) \propto p_{\mathbf{y}}(\mathbf{y}_k | \mathbf{x}') p_{\mathbf{x}}(\mathbf{x}', t_{k-}). \quad (22)$$

From this, it is obvious that, when combining a GGDAF prior $p_{\mathbf{x}}(\mathbf{x}', t_{k-})$ with a measurement with Gaussian uncertainty $p_{\mathbf{y}}(\mathbf{y}_k | \mathbf{x}')$ (as is commonly the case), the posterior $p_{\mathbf{x}}(\mathbf{x}', t_{k+})$ is a GGDAF; therefore the GGDAF family of distributions are conjugate priors.

5. Numerical probability distributions

Now consider an arbitrary relaxation shape $\bar{p}(\mathbf{x})$, obtained experimentally and appropriately smoothed. For instance, when modeling the motions of an apex predator within its observed territory, one might collect a discrete set of measurements representing the coordinates of the animal. Using this set, one would like to approximately

define the ROI that represents the non-evasive statistics of said animal. The ROI could then be used to develop a steady-state $\bar{p}(\mathbf{x})$, from which the relaxation advection field $\bar{\mathbf{v}}(\mathbf{x})$ could be approximated via a finite difference approximation of (4). To focus this discussion, consider the Fokker-Planck PDE discretized on a uniform Cartesian 2D mesh (with grid width Δx and Δy), noticing that higher-dimensional cases follow as an obvious extension. Assuming homogeneous Dirichlet boundary conditions, the advection field $\bar{\mathbf{v}}_{i,j}$ representing the advection at coordinate (i, j) can be written as

$$\bar{\mathbf{v}}_{i,j} = \frac{D}{\bar{p}_{i,j}} \left[\frac{\bar{p}_{i+1,j} - \bar{p}_{i-1,j}}{2\Delta x} \quad \frac{\bar{p}_{i,j+1} - \bar{p}_{i,j-1}}{2\Delta y} \right]^T. \quad (23)$$

An important consideration when approximating $\bar{\mathbf{v}}_{i,j}$ is the singularity that can occur at $\bar{p}_{i,j} = 0$. We can circumvent this issue by assuming that p is a sufficiently smooth function, or infinitely differentiable, thus $\bar{p}_{i,j} = 0$ only at the boundaries. Because of our homogeneous Dirichlet boundary conditions, $\bar{\mathbf{v}}_{i,j}$ at the boundaries is negligible. Using the approximated $\bar{\mathbf{v}}_{i,j}$, an arbitrary PDF can be relaxed in the same manner as in the analytical examples. However, generating a $\bar{p}(\mathbf{x})$ from a dataset is nontrivial; the two different methods that this paper focuses on are the α -convex hull approximation and kernel density estimation.

5.1. α -convex hull approximation

The α -convex hull [38] is a generalization of the convex hull; where the convex hull assumes convexity of the estimator region Ω , the α -convex hull assumes α -convexity, a more flexible constraint that enlarges the family of sets that may be estimated. For each $\alpha \in \mathbb{R}$, a “generalized disk” of radius $1/\alpha$ is defined as follows: if $\alpha = 0$, the disk is a closed halfplane (resulting in an ordinary convex hull), if $\alpha > 0$ the disk is closed with a radius of $1/\alpha$, and if $\alpha < 0$ the disk is the closure of the complement of a disk of radius $-1/\alpha$. Using this definition of a disk, the α -convex hull, or ROI Ω , of a dataset S may be defined as the intersection of all closed disks that contain all points of S .

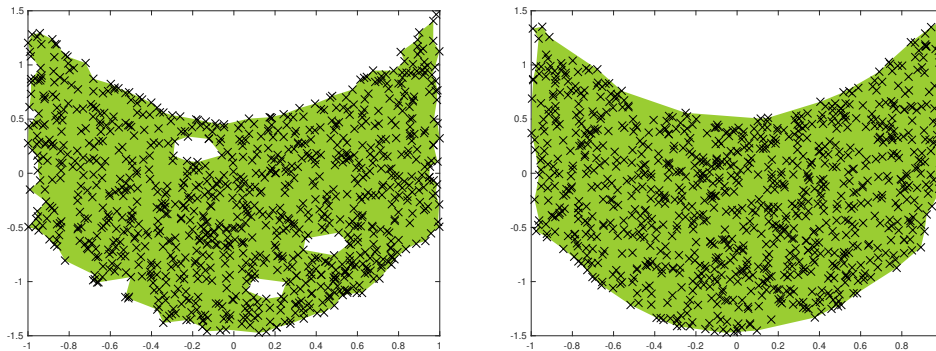


Figure 9. Estimation of the boundary of an unknown ROI Ω (taken in this as the space between two parabolas) based on 1000 experimentally obtained datapoints randomly distributed over Ω , using the alpha shape algorithm, as implemented in Matlab, using (*left*) an alpha radius (for adjusting the tightness of the fit around the points) of 0.1, and (*right*) an alpha radius of 0.4. This sophisticated algorithm, which is built on a Delaunay triangulation of the available datapoints, is a convenient generalization of the more traditional convex hull algorithm.

Converting Ω into a d -dimensional PDF requires some smoothing via diffusion. Ω

is first converted to a uniform distribution $p_\Omega(\mathbf{x})$, defined as

$$p_\Omega(\mathbf{x}) = \frac{1}{C} \begin{cases} 1, & \mathbf{x} \in \Omega \\ 0, & \text{else} \end{cases}, \quad (24)$$

where C is such that $\int_\Omega p_\Omega(\mathbf{x}) d\mathbf{x} = 1$. To smooth the edges, the PDF is CT evolved with the diffusion-only finite difference stencil $\mathcal{A}(D)$. Consider the uniform Cartesian 2D mesh finite difference approximation of (3) when $\bar{\mathbf{v}}(\mathbf{x}) = 0$ (assuming $D = \lambda I$)

$$\frac{p_{i,j}^{n+1} - p_{i,j}^n}{\Delta t} = \lambda \left(\frac{p_{i+1,j}^n - 2p_{i,j}^n + p_{i-1,j}^n}{\Delta x^2} + \frac{p_{i,j+1}^n - 2p_{i,j}^n + p_{i,j-1}^n}{\Delta y^2} \right) \quad (25)$$

The finite difference stencil $\mathcal{A}(D)$ representing (25) is

$$\mathcal{A}(D) = \begin{bmatrix} 0 & \frac{\lambda}{\Delta y^2} & 0 \\ \frac{\lambda}{\Delta x^2} & -\left(\frac{2\lambda}{\Delta x^2} + \frac{2\lambda}{\Delta y^2}\right) & \frac{\lambda}{\Delta x^2} \\ 0 & \frac{\lambda}{\Delta y^2} & 0 \end{bmatrix}, \quad (26)$$

The stencil is used to form the block tridiagonal Toeplitz matrix M (as demonstrated in Chapter 1 of [36]) which is used for propagating $p_\Omega(\mathbf{x})$ until it is sufficiently smoothed. Using the explicit Euler method, this is written as

$$p_\Omega^{n+1}(\mathbf{x}) = p_\Omega^n(\mathbf{x}) + \Delta t \left(M p_\Omega^n(\mathbf{x}) \right). \quad (27)$$

The number of timesteps of propagation necessary for sufficient smoothing is relative to the problem and can be determined empirically. The final result is a smoothed, d -dimensional $\bar{p}(\mathbf{x})$ representing the ROI Ω . From here, (23) can be used to calculate $\bar{\mathbf{v}}(\mathbf{x})$ in a discretized manner. Application of this smoothing process with 20 timesteps of smoothing, with $\Delta t = 0.001$ and $\lambda = 1$ on the dataset from Figure 9, along with the corresponding $\bar{\mathbf{v}}(\mathbf{x})$, is shown in Figure 10.




Figure 10. (left) $\bar{p}(\mathbf{x})$ and (right) $\bar{\phi}(\mathbf{x} | \lambda = 1)$ of numerically defined ROI Ω , approximated via α -convex hull estimation of sample dataset from Figure 9.

5.2. Kernel density estimation approximation

Kernel density estimation (KDE) [39–42] is a non-parametric technique that approximates the PDF of a random variable \mathbf{x} given a set of k realizations of said random variable \mathbf{x}_i , $i = 1, \dots, k$. KDE generates the underlying PDF by placing a kernel function $K(\cdot)$ scaled by a bandwidth h at each realization. It then sums up the kernels and normalizes to form a density function. Explicitly stated, the d -dimensional PDF $p_h(\mathbf{x})$ generated by KDE is written as

$$p_h(\mathbf{x}) = \frac{1}{k h^n} \sum_{i=1}^k K\left(\frac{\mathbf{x} - \mathbf{x}_i}{h}\right). \quad (28)$$

$K(\cdot)$ is typically a symmetric function that peaks at the realization value \mathbf{x}_i , and decays to zero moving away from it. Commonly used kernel functions are the Gaussian, box, triangle, and Epanechnikov functions. The bandwidth h determines the smoothness of $p_h(x)$ and is dependent on the kernel function. A larger bandwidth will result in a smoother density, while a smaller bandwidth may capture more local behavior. Generally, KDE is utilized when an excess of data exists and said data indicates areas of bias within Ω . For instance, consider an apex predator with a territory similar to that of Figure 9, with the key difference being that it sleeps near the edge of its territory. In this case, the most likely area within the ROI for the predator to be is no longer at the center. In another scenario, there exists an uninhabitable area within the predator’s ROI, like a lake. As the predator has no means of entering this subregion, the probability should be negligible here. For both cases, this information would be lost had Ω been approximated via the α -convex hull (depending on the value of α , for the latter case).




Figure 11. Two distinct types of territories, both containing the nominal dataset from Figure 9; one (*left*) includes additional points sampled from the red ellipse, representing a den, and the other (*right*) removes all points from the green ellipse, representing a lake.

Following the constraints from the datasets from Figure 11, denoted “den dataset” and “lake dataset” for left and right datasets respectively, and increasing the number of points to $k = 10,000$ for additional smoothing, the KDE algorithm in Matlab is utilized, with the default Gaussian kernel function and bandwidth $h = 0.1$. With 5 timesteps of smoothing at $\Delta t = 0.001$ and $\lambda = 1$, $\bar{p}(\mathbf{x})$ and the corresponding $\bar{\mathbf{v}}(\mathbf{x})$ are approximated, shown in Figure 12.

Determining when to use the α -convex hull approximation versus KDE depends

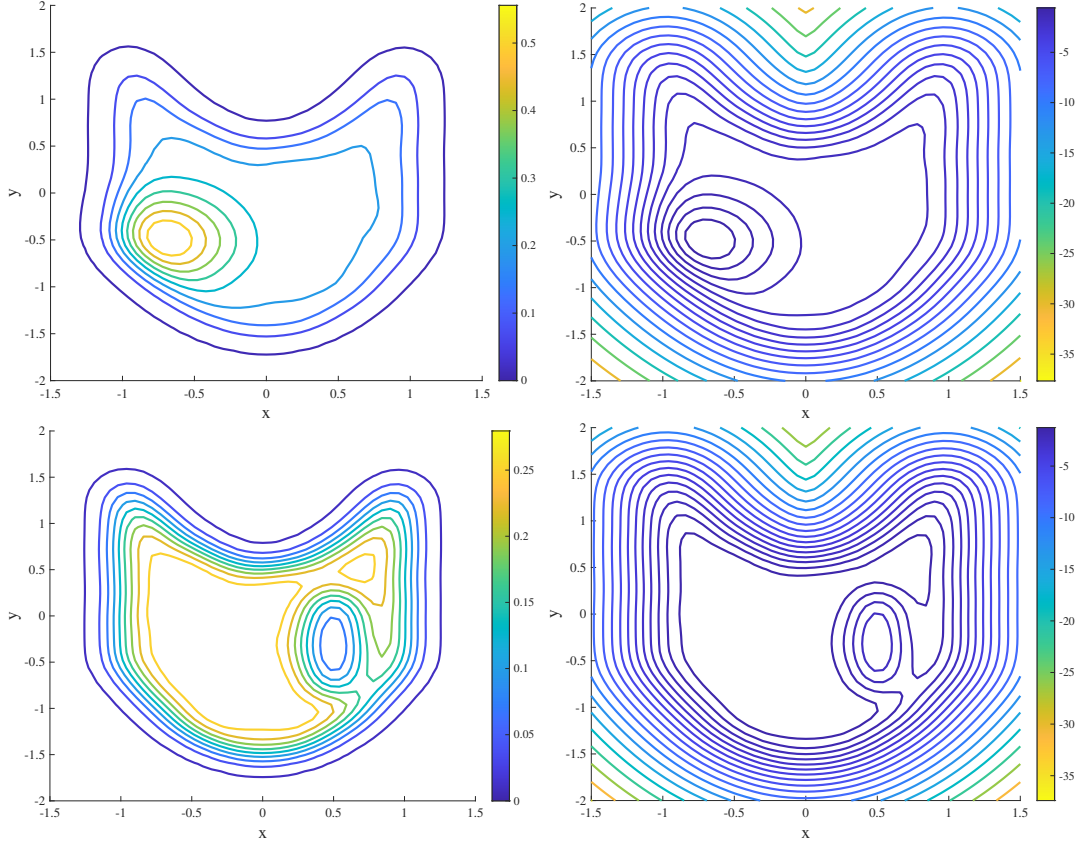


Figure 12. For the den dataset: (*top left*) $\bar{p}(\mathbf{x})$ and (*top right*) $\bar{\phi}(\mathbf{x} | \lambda = 1)$ of numerically defined ROI Ω . For the lake dataset: (*bottom left*) $\bar{p}(\mathbf{x})$ and (*bottom right*) $\bar{\phi}(\mathbf{x} | \lambda = 1)$ of numerically defined ROI Ω , both determined via KDE.

on the characteristics of the dataset. When data is sparse, using the α -convex hull approximation may more accurately portray the PDF of the non-evasive target compared to KDE. However, the α -convex hull approximation assumes that each point within Ω is an equally likely location of the target. By smoothing the ROI via diffusion, this creates a $\bar{p}(\mathbf{x})$ with highest probability near the center of Ω and decaying probability moving outward. Although this may be a realistic scenario, it is not always the case. When data is plentiful, KDE can capture areas of focus otherwise ignored via the α -convex hull approximation. If data is sparse, the bandwidth of KDE can cause smoothing that removes important local variations that would've been well represented with the α -convex hull approximation.

6. Application of Relaxation Advection

Utilizing the framework previously developed, we now demonstrate the application of the derived relaxation advection field with diffusion, the application of which drives an arbitrary PDF (taken to be a uniform distribution for this section) to a pre-determined, either analytic or numerical, steady-state $\bar{p}(\mathbf{x})$.

6.1. CT evolving using advection-diffusion finite stencil

Once $\bar{\mathbf{v}}(\mathbf{x})$ is calculated, either analytically or numerically, it can be inserted into (3) to drive an arbitrary PDF to the relaxation shape $\bar{p}(\mathbf{x})$. Similar to the process of smoothing with diffusion, a finite difference stencil is utilized, this time with advection included $\mathcal{A}(\mathbf{v}, D)$. Consider the uniform Cartesian 2D mesh finite difference approximation of (3) with the relaxation advection field $\bar{\mathbf{v}}(\mathbf{x})$ applied (again with isotropic diffusion $D = \lambda I$)

$$\begin{aligned} \frac{p_{i,j}^{n+1} - p_{i,j}^n}{\Delta t} + p_{i,j}^n \left(\frac{\bar{v}_{i+1,j}^x - \bar{v}_{i-1,j}^x}{2\Delta x} + \frac{\bar{v}_{i,j+1}^y - \bar{v}_{i,j-1}^y}{2\Delta y} \right) + \left(\bar{v}_{i,j}^x \frac{p_{i+1,j}^n - p_{i-1,j}^n}{2\Delta x} + \bar{v}_{i,j}^y \frac{p_{i,j+1}^n - p_{i,j-1}^n}{2\Delta y} \right) \\ = \lambda \left(\frac{p_{i+1,j}^n - 2p_{i,j}^n + p_{i-1,j}^n}{\Delta x^2} + \frac{p_{i,j+1}^n - 2p_{i,j}^n + p_{i,j-1}^n}{\Delta y^2} \right) \end{aligned} \quad (29)$$

where $\bar{\mathbf{v}}_{i,j} = \begin{bmatrix} \bar{v}_{i,j}^x \\ \bar{v}_{i,j}^y \end{bmatrix}$. Thus, the finite difference stencil $\mathcal{A}(\bar{\mathbf{v}}, D)$, utilized for all of the following examples, is

$$\mathcal{A}(\bar{\mathbf{v}}, D) = \begin{bmatrix} 0 & \frac{\lambda}{\Delta y^2} - \frac{\bar{v}_{i,j}^y}{2\Delta y} & 0 \\ \frac{\lambda}{\Delta x^2} + \frac{\bar{v}_{i,j}^x}{2\Delta x} - \left(\frac{2\lambda}{\Delta x^2} + \frac{2\lambda}{\Delta y^2} \right) - \left(\frac{\bar{v}_{i+1,j}^x - \bar{v}_{i-1,j}^x}{2\Delta x} + \frac{\bar{v}_{i,j+1}^y - \bar{v}_{i,j-1}^y}{2\Delta y} \right) & \frac{\lambda}{\Delta x^2} - \frac{\bar{v}_{i,j}^x}{2\Delta x} & 0 \\ 0 & \frac{\lambda}{\Delta y^2} + \frac{\bar{v}_{i,j}^y}{2\Delta y} & 0 \end{bmatrix}. \quad (30)$$

The stencil is used to CT update an arbitrary PDF in the same manner as aforementioned. To avoid bias, the initial PDF to be relaxed is a uniform distribution over Ω .

6.2. Convergence to relaxation shape

To determine if the driven PDF $p(\mathbf{x}, t)$ has converged to the relaxation shape $\bar{p}(\mathbf{x})$, the absolute relative difference is computed at each grid point, then summed up to compute ϵ . On a uniform Cartesian 2D mesh square grid of size N , ϵ can be written as

$$\epsilon = \sum_{i=1}^N \sum_{j=1}^N |p_{i,j} - \bar{p}_{i,j}|. \quad (31)$$

Convergence was assumed when $\Delta\epsilon = \epsilon(t_{k+1}) - \epsilon(t_k) < 0.01$, or the summation of the absolute relative difference at timestep t_{k+1} minus the summation of the absolute relative difference at timestep t_k . In both analytical examples provided, we use parameters $\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} 1 & 0.4 \\ 0.4 & 1 \end{bmatrix}$. The parameters for the frames from the simulations can be found in Table 1.

7. Probabilistic search

Probabilistic search algorithms employ statistical methods that find optimal solutions in solution-space [32]. For target searching specifically, these algorithms utilize the information provided by the PDF of a target to control the trajectories of the search




Figure 13. A uniform distribution relaxed to a GD $\bar{p}(\mathbf{x}|\boldsymbol{\mu},\Sigma)$ via the analytical relaxation advection $\bar{\mathbf{v}}(\mathbf{x}|\boldsymbol{\mu},\Sigma)$ at timestep (left) 50, (middle) 100, and (right) convergence.




Figure 14. A uniform distribution relaxed to a GGD $\bar{p}(\mathbf{x}|\boldsymbol{\mu},\Sigma,\beta=2)$ via the analytical relaxation advection $\bar{\mathbf{v}}(\mathbf{x}|\boldsymbol{\mu},\Sigma,\beta=2)$ at timestep (left) 50, (middle) 100, and (right) convergence.




Figure 15. A uniform distribution relaxed to an α -convex hull approximation $\bar{p}(\mathbf{x})$ via the numerical relaxation advection $\bar{\mathbf{v}}(\mathbf{x})$ at timestep (left) 100, (middle) 400, and (right) convergence, approximated based on the dataset from Figure 9.




Figure 16. A uniform distribution relaxed to a KDE approximation $\bar{p}(\mathbf{x})$ via the numerical relaxation advection $\bar{\mathbf{v}}(\mathbf{x})$ at timestep (left) 100, (middle) 400, and (right) convergence, approximated based on the den dataset from Figure 11.




Figure 17. A uniform distribution relaxed to a KDE approximation $\bar{p}(\mathbf{x})$ via the numerical relaxation advection $\bar{\mathbf{v}}(\mathbf{x})$ at timestep (*left*) 100, (*middle*) 400, and (*right*) convergence, approximated based on the lake dataset from Figure 11.

vehicles examining the ROI, with the objective of minimizing the search time required to detect said target. As the dynamics of the target may vary, (a non-evasive target versus an evasive target, for example), so too may the CT evolution of the PDF, governed by (3). We provide a class of probabilistic search algorithm that utilizes uninformed, periodic search vehicle trajectories as a means for displaying the capabilities of non-evasive and evasive target searching utilizing the relaxation advection $\bar{\mathbf{v}}(\mathbf{x})$ previously derived. This work is meant to act as a straightforward starting point for a trajectory optimization method, with search trajectories that are adaptive and informed by the PDF of the target.

7.1. Non-evasive targets

For the probabilistic search of a non-evasive target, we assume $D(\mathbf{x})$ and $\mathbf{v}(\mathbf{x})$ are time-invariant, i.e. the target is statistically stationary. Then, we use the relaxation advection field $\bar{\mathbf{v}}(\mathbf{x})$ from (4) and homogenous isotropic diffusion $D = \lambda_0 I$ to represent the non-evasive target. We introduce $m = 1, \dots, M$ search vehicles moving in periodic orbits about the ROI. The search vehicles have sensors that form a limited field of view of the ROI, modeled in the form of a normal distribution $p_m(\mathbf{x}, t) \sim \mathcal{N}(\mathbf{q}_m(t), \sigma_m^2)$ where $\mathbf{q}_m(t)$ is the position and σ_m^2 is the effective width of the field of view of the m^{th} search vehicle, respectively. This density function acts as an external forcing term in (3) that “mows” down the probability in its vicinity, assuming that it fails to locate the target. In this model, where we account for observations, the Fokker-Planck PDE becomes

$$\frac{\partial p(\mathbf{x}, t)}{\partial t} - \nabla_{\mathbf{x}} \cdot \left[\lambda_0 \nabla_{\mathbf{x}} p(\mathbf{x}, t) - \bar{\mathbf{v}}(\mathbf{x}) p(\mathbf{x}, t) \right] = -p(\mathbf{x}, t) \left(b(\mathbf{x}, t) - \gamma(t) \right), \quad (32a)$$

$$\text{where } b(\mathbf{x}, t) = \sum_{m=1}^M a_m \exp \left\{ -\frac{\|\mathbf{x} - \mathbf{q}_m(t)\|^2}{2\sigma_m^2} \right\} \quad \text{and} \quad \gamma(t) = \int_{\Omega} p(\mathbf{x}, t) b(\mathbf{x}, t) d\mathbf{x}. \quad (32b)$$

The normalization factor $\gamma(t)$ ensures that $\int_{\Omega} p(\mathbf{x}, t) = 1$ for all t , and a_m is the acuity of the sensor of the m^{th} search vehicle (this model assumes the sensor’s range decays exponentially as a function of distance, but other distributions may be used). For the scope of this paper, the trajectory of the search vehicles $\mathbf{q}_m(t)$ are fixed and are designed to best cover the known ROI. Examples of such trajectories are concentric circles, Cassini ovals [43], and offset lemniscates [44]. Note that larger values of λ lead to quicker flow of probability, representing a faster-moving target, while smaller values of λ , representing a slower-moving target, may lead to buildups of probability in small

subregions. Both behaviors are demonstrated in Figure 18.

7.2. Evasive targets

For an evasive target, we no longer assume the statistics are time-invariant, thus $D(\mathbf{x}, t)$ and $\mathbf{v}(\mathbf{x}, t)$ are both functions of space and time. Consider a skittish deer, a target that is both aware of the search vehicles and actively evading them, with increasing “agitation”, or random fluctuation, and “evasiveness”, or advection away from, as the search vehicles come closer. In this case, we model the evasive, isotropic diffusion $D(\mathbf{x}, t)$ as the relaxation diffusion tensor $\lambda_0 I$ plus the superposition of GGDs centered at the locations of the M search vehicles, and the evasive advection $\tilde{\mathbf{v}}(\mathbf{x}, t)$ as the relaxation advection $\bar{\mathbf{v}}(\mathbf{x})$ plus the superposition of the negative gradients of GGDs centered at the locations of the M search vehicles, explicitly stated as

$$D(\mathbf{x}, t) = \lambda(\mathbf{x}, t)I, \quad \text{where} \quad \lambda(\mathbf{x}, t) = \lambda_0 + \psi \sum_{m=1}^M p_m(\mathbf{x} | \mathbf{q}_m(t), \Sigma_m), \quad (33a)$$

$$\text{and} \quad \tilde{\mathbf{v}}(\mathbf{x}, t) = \bar{\mathbf{v}}(\mathbf{x}) - \delta \sum_{m=1}^M \nabla_{\mathbf{x}} p_m(\mathbf{x} | \mathbf{q}_m(t), \Sigma_m), \quad \text{where} \quad (33b)$$

$$p_m(\mathbf{x} | \mathbf{q}_m(t), \Sigma_m) = A_{\beta, d} \exp \left\{ - \left[B_{\beta, d} [\mathbf{x} - \mathbf{q}_m(t)]^T \Sigma_m^{-1} [\mathbf{x} - \mathbf{q}_m(t)] \right]^\beta \right\}, \quad (33c)$$

$\psi, \delta \geq 0$ are the agitation and evasiveness parameters of the target, respectively. In the case of inhomogenous, time-varying advection and diffusion, the Fokker-Planck PDE becomes

$$\begin{aligned} \frac{\partial p(\mathbf{x}, t)}{\partial t} - \nabla_{\mathbf{x}} \cdot \left[\lambda(\mathbf{x}, t) \nabla_{\mathbf{x}} p(\mathbf{x}, t) + p(\mathbf{x}, t) [\nabla_{\mathbf{x}} \cdot \lambda(\mathbf{x}, t) I] - \tilde{\mathbf{v}}(\mathbf{x}, t) p(\mathbf{x}, t) \right] \\ = -p(\mathbf{x}, t) \left(b(\mathbf{x}, t) - \gamma(t) \right). \end{aligned} \quad (34)$$

The prior distribution used as an initialization point for the target search is the relaxation PDF $\bar{p}(\mathbf{x})$. For evasive targets, we utilize cooperative hunting, a strategy used by pack hunters that drives the prey towards the “kill zone” (or generally, the identification zone), considering their avoidance behavior, commonly used by dolphins [8]. Using rotating trajectories from the Cassini oval family, the search vehicles begin at the outskirts of the ROI, directing the target inwards, then converge simultaneously from all directions, trapping the target in the center, as demonstrated in Figure 19.

7.3. Results

To demonstrate the probabilistic search on non-evasive and evasive targets, we discretize the dynamics from (32a) and (34) on an $N \times N$ Cartesian grid of width Δx and Δy and time-march $\bar{p}(\mathbf{x}, t)$ using an RK4 scheme and a finite difference stencil M aforementioned. In Figure 18, we demonstrate three different, non-evasive ROI’s:

- (1) $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta = 1)$ where $\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} 3 & 0 \\ 0 & 3 \end{bmatrix}$ with concentric circular orbits
- (2) $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \Sigma, \beta = 2)$ where $\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} 2 & 0.8 \\ 0.8 & 2 \end{bmatrix}$ with Cassini ovals




Figure 18. Frames from three different probabilistic search simulations of non-evasive targets with dynamics from (32a). The solid lines represent the fixed orbits of the search vehicles and the points are the search vehicles' positions. Parameters for simulations can be found in Table 2, rows 1-3.

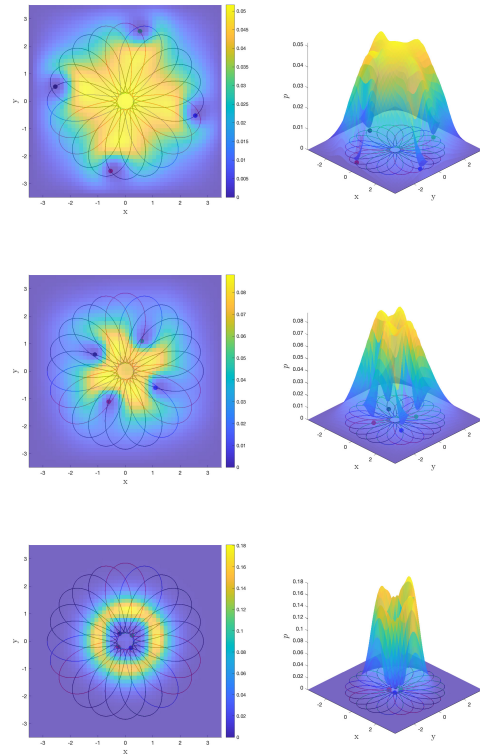


Figure 19. Three subsequent frames from the same probabilistic search simulation of an evasive target with dynamics from (34). The solid lines represent the fixed orbits of the search vehicles and the points are the search vehicles' positions. Parameters for simulations can be found in Table 2, row 4.

- (3) An α -convex hull approximation from the dataset from Figure 10 with offset, scaled lemniscates

We demonstrate an evasive target search in Figure 19, utilizing a set of rotating shapes from the Cassini oval family. From the subsequent frames, we illustrate how the herding vehicles increase the probability density at the center of the ROI, leading to an increased likelihood of target discovery. The full list of parameters for the simulations can be found in Table 2 and the location of the simulation animations can be found in the Appendix.

8. Conclusion

This paper effectively demonstrates the CT framework for a probabilistic search algorithm that evolves the PDF $p(\mathbf{x}, t)$ of a stochastically-moving, statistically stationary target governed by the steady-state PDF $\bar{p}(\mathbf{x})$ summarizing the ROI $\Omega \in \mathbb{R}^d$ of the target, which may be defined analytically or experimentally. Using the Fokker-Planck PDE, we solve for the relaxation advection equation $\bar{\mathbf{v}}(\mathbf{x})$, shown to be only a function of the diffusion tensor $D(\mathbf{x})$ and $\bar{p}(\mathbf{x})$, which, when applied with $D(\mathbf{x})$, drives an arbitrary $p(\mathbf{x}, t)$ to steady-state $\bar{p}(\mathbf{x})$. In the analytical regime, for the Gaussian

family of distributions, we use the closed-form $\bar{p}(\mathbf{x})$ to solve for a closed-form solution to $\bar{\mathbf{v}}(\mathbf{x})$. In the experimental regime, we demonstrate how a set of observations may be used to numerically solve for $\bar{p}(\mathbf{x})$, either using the α -convex hull approximation or the kernel density estimation; the decision between the two methods was also discussed at length. For either numerical technique, $\bar{\mathbf{v}}(\mathbf{x})$ may be derived numerically from the steady-state solution of the Fokker-Planck PDE. The efficacy of using either the analytically- or numerically-defined relaxation advection was demonstrated by driving uniform distributions to steady-state.

Having validated $\bar{\mathbf{v}}(\mathbf{x})$, we set up a framework for how it may be used to perform a probabilistic search. First, we consider the behavior of the target. For a non-evasive target, the presence of search vehicles will have no impact on the motion, thus the advection in the Fokker-Planck PDE is just the relaxation advection. To the evolution equation of the PDF, we introduce an external forcing term $b(\mathbf{x}, t)$, representative of search vehicle observations that decrease the probability of discovery in their vicinity. To ensure unity of the PDF, we include a normalization term $\gamma(t)$. We then RK4 time-march the initially steady-state PDF subject to the listed forces and demonstrate how the PDF changes over time, with the assumption that the observations do not result in discovery. The search trajectories $\mathbf{q}_m(t)$ for the non-evasive target are uninformed, periodic orbits of varying shape. For an evasive target, we add to the relaxation advection a tendency to move away from, or evade, the search vehicles. To the isotropic diffusion we add a tendency to become more agitated, or random, as the search vehicles get closer. Including $b(\mathbf{x}, t)$ and $\gamma(t)$, we again RK4 time-march the initially steady-state PDF with the assumption that the observations do not result in discovery. The search trajectories $\mathbf{q}_m(t)$ for the evasive target, while still uninformed with regards to $p(\mathbf{x}, t)$, are representative of a coordinate herding technique that is meant to drive the target to a central identification zone, where discovery is more likely.

As opposed to other search techniques that stochastically evolve the state of a target, the technique we propose evolves the full PDF, which may provide more information to search vehicle optimization algorithms, improving the standards for search and rescue procedures, resulting in targets being located faster. In our demonstrations, the search trajectories are uninformed periodic orbits meant to cover the ROI and do not take advantage of the real-time information provided by $p(\mathbf{x}, t)$. Future work will look to change these search trajectories, optimizing and directing $\mathbf{q}_m(t)$ to regions of high probability within the ROI for fast target procurement.

As a medium for testing the probabilistic search technique, we derive a new formulation for the multivariate Generalized Gaussian Distribution, an extension of the Gaussian Distribution raised to β . The formulation we present includes a change to the standard definition of the multivariate GGD to ensure that the second central moment, or the covariance, is the matrix Σ , as is the case for the GD. We theorize that this new formulation of the GGD may be advantageous for parameter estimation algorithms in the field of mixture modeling, particularly for determining the closed-form solutions to the parameter update steps. Finally, to enhance the flexibility of the GGD for more precisely fitting data, we introduce the Generalized Gaussian Distribution with Anisotropic Flatness (GGDAF), a GGD with shape parameters that may differ along the eigenvectors of Σ . We formulate the 2D and 3D GGDAF such that the distribution simplifies to the GGD when the shaping parameters are equal and displays anisotropic flatness otherwise. The presented formulation of the GGDAF is only suited for integer values of β , but future work will aim to address this limitation.

Appendix

The code to recreate all of the results provided in this paper can be found at <https://github.com/bhanson10/Prob-Search>. All animations can be found at <https://bhanson10.github.io/PS.html> using password “FCCR”.

Table 1. Relaxation advection simulation parameters

	Frame	1			2			3 (convergence)		
	Parameter	iter	t	ϵ	iter	t	ϵ	iter	t	ϵ
Analytical	Gaussian	50	0.25	41.66	100	0.5	20.53	246	1.23	2.21
	GGD, $\beta = 2$	50	0.15	48.87	100	0.3	26.70	400	1.2	1.50
Numerical	α -convex hull	100	0.1	93.68	400	0.4	23.17	1021	1.02	4.94
	KDE - den	100	0.1	104.26	400	0.4	37.71	1301	1.30	6.79
	KDE - lake	100	0.1	86.09	400	0.4	23.71	1008	1.01	7.27

Table 2. Probabilistic search simulation parameters

#	Target characteristics				Search vehicle orbit characteristics	
	Process type	Prior shape	Parameters	λ	Orbit shape	Parameters
1	Non-evasive	Gaussian	$\mu = \begin{bmatrix} 0 \\ 0 \\ 3 \\ 0 \end{bmatrix}$ $\Sigma = \begin{bmatrix} 0 & & & \\ & 3 & & \\ & & 0 & \\ & & & 3 \end{bmatrix}$	1	Concentric circles	radii = [0.5, 1.5, 2.5] $\theta_0 = [0, \frac{2\pi}{3}, \frac{4\pi}{3}]$ $\dot{\theta} = 8$ $\sigma_m = 0.25$ $a_m = 1e4$
2		GGD	$\beta = 2$ $\mu = \begin{bmatrix} 0 \\ 0 \\ 2 \\ 0.8 \end{bmatrix}$ $\Sigma = \begin{bmatrix} 2 & & & \\ & 0.8 & & \\ & & 0.8 & \\ & & & 2 \end{bmatrix}$	0.5	Cassini ovals	$A = 1.2$ $B = [1.22, 1.4, 1.7]$ $f = 1.41$ $\theta_0 = [0, \frac{2\pi}{3}, \frac{4\pi}{3}]$ $\dot{\theta} = 8$ $\sigma_m = 0.25$ $a_m = 4e3$
3		α -convex hull	Figure 9 dataset	0.25	Offset lemniscates	$f = 1.4$ $\theta_0 = [0, \frac{2\pi}{3}, \frac{4\pi}{3}]$ $\dot{\theta} = 8$ $\sigma_m = 0.1$ $a_m = 8e3$
4	Evasive	Gaussian	$\mu = \begin{bmatrix} 0 \\ 0 \\ 3 \\ 0 \end{bmatrix}$ $\Sigma = \begin{bmatrix} 0 & & & \\ & 3 & & \\ & & 0 & \\ & & & 3 \end{bmatrix}$ $\delta = 8$ $\psi = 0.1$	2	Rotating Cassini ovals	$A = 2$ $B = [2.02, 2.021, 2.022, 2.023]$ $v = 0.4$ $\theta_0 = [0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}]$ $\dot{\theta} = 1$ $\sigma_m = 0.1$ $a_m = 4e2$ $s = 0.6$

Acknowledgments

The authors gratefully acknowledge Prof. Paulo Luchini and Prof. Boris Kramer for insightful discussion related to this work.

References

- [1] Chung TH. On probabilistic search decisions under searcher motion constraints ; 2010.

- [2] Chung TH, Burdick JW. Analysis of search decision making using probabilistic search strategies. *IEEE Transactions on Robotics*. 2012;28(1):132–144.
- [3] Noori N, Renzaglia A, Hook JV, et al. Constrained probabilistic search for a one-dimensional random walker. *IEEE Transactions on Robotics*. 2016;32(2):261–274.
- [4] Washburn AR. Branch and bound methods for a search problem. *Naval Res Logistics*. 1998;45(3):243–257.
- [5] Hollinger G, Singh S, Djughash J, et al. Efficient multi-robot search for a moving target. *Int J Robot Res*. 2009;28(2):201–219.
- [6] Chew MC. Optimal stopping in a discrete search problem. *Operations Research*. 1973; 21(3):741–747.
- [7] Ross SM. A problem in optimal search and stop. *Operations Research*. 1969;17(6):984–992.
- [8] Klopfer P. Habitats and territories: a study of the use of space by animals. *Habitats And Territories A Study Of The Use Of Space By Animals*. 1969;.
- [9] Van Moorter B, Visscher D, Benhamou S, et al. Memory keeps you at home: a mechanistic model for home range emergence. *Oikos*. 2009;118(5):641–652.
- [10] Powell R. Animal home ranges and territories and home range estimators ; 2000.
- [11] Benhamou S. Of scales and stationarity in animal movements. *Ecology Letters*. 2014; 17(3):261–272.
- [12] Patin R, Etienne MP, Lebarbier E, et al. Identifying stationary phases in multivariate time series for highlighting behavioural modes and home range settlements. *Journal of Animal Ecology*. 2020;89(1):44–56.
- [13] Kaufman E. New details in F-35 'mishap' as mystery remains about how jet was lost ; 2023. ABC News.
- [14] Live air trafic www.flightradar24.com/about ; 2023. FlightRadar24.
- [15] Gauss CF. *Theoria combinationis observationum erroribus minimis obnoxiae*. Henricus Dieterich; 1823.
- [16] Subbotin MT. On the law of frequency of error. *Mathematical Collection*. 1923;31(2):296–301.
- [17] Taguchi T. On a generalization of gaussian distribution. *Annals of the Institute of Statistical Mathematics/edited by the Institute of Statistical Mathematics*. 1978;30(2):p211–242.
- [18] Coban MZ, Mersereau RM. Adaptive subband video coding using bivariate generalized gaussian distribution model. In: 1996 IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings; Vol. 4; IEEE; 1996. p. 1990–1993.
- [19] Gómez E, Gomez-Viilegas M, Marín JM. A multivariate generalization of the power exponential family of distributions. *Communications in Statistics-Theory and Methods*. 1998; 27(3):589–600.
- [20] Verdoolaege G, Scheunders P. On the geometry of multivariate generalized gaussian models. *Journal of mathematical imaging and vision*. 2012;43:180–193.
- [21] Zhang T, Wiesel A, Greco MS. Multivariate generalized gaussian distribution: Convexity and graphical models. *IEEE Transactions on Signal Processing*. 2013;61(16):4141–4148.
- [22] Bouguila N, Fan W. *Mixture models and applications*. Springer; 2020.
- [23] Sorenson HW, Alspach DL. Recursive bayesian estimation using gaussian sums. *Automatica*. 1971;7(4):465–479.
- [24] Alspach D, Sorenson H. Nonlinear bayesian estimation using gaussian sum approximations. *IEEE transactions on automatic control*. 1972;17(4):439–448.
- [25] Yun S, Zanetti R. Sequential monte carlo filtering with gaussian mixture sampling. *Journal of Guidance, Control, and Dynamics*. 2019;42(9):2069–2077.
- [26] Yun S, Zanetti R, Jones BA. Kernel-based ensemble gaussian mixture filtering for orbit determination with sparse data. *Advances in Space Research*. 2022;69(12):4179–4197.
- [27] Mohamed OMM, Jaïdane-Saïdane M. Generalized gaussian mixture model. In: 2009 17th European signal processing conference; IEEE; 2009. p. 2273–2277.
- [28] Nguyen TM, Wu QJ, Zhang H. Bounded generalized gaussian mixture model. *Pattern Recognition*. 2014;47(9):3132–3142.

- [29] Lee TW, Lewicki MS. The generalized gaussian mixture model using ica. In: International Workshop on Independent Component Analysis (ICA'00); 2000. p. 239–244.
- [30] Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*. 1977; 39(1):1–22.
- [31] McLachlan GJ, Krishnan T. *The em algorithm and extensions*. John Wiley & Sons; 2007.
- [32] Stone L. *Theory of optimal search*. 2nd ed. Military application section, Operations Research Society of America; 1992.
- [33] Fokker AD. Die mittlere energie rotierender elektrischer dipole im strahlungsfeld. *Ann Phys*. 1914;348:810–820.
- [34] Planck M. Über einen satz der statistischen dynamik und seine erweiterung in der quantentheorie. Reimer; 1917. *Sitzungsberichte der Königlich-Preussischen Akademie der Wissenschaften zu Berlin*.
- [35] Guerrero JS, Pimentel LCG, Skaggs TH, et al. Analytical solution of the advection–diffusion transport equation using a change-of-variable and integral transform technique. *International Journal of Heat and Mass Transfer*. 2009;52:3297–3304.
- [36] Bewley T. *Numerical renaissance: simulation, optimization, control*. Renaissance Press; 2018.
- [37] Bayes T. An essay towards solving a problem in the doctrine of chances. *Phil Trans of the Royal Soc of London*. 1763;53:370–418.
- [38] Edelsbrunner H, Kirkpatrick DG, Seidel R. On the shape of a set of points in the plane. *IEEE Transactions on Information Theory*. 1983;29 (4):551–559.
- [39] Bowman AW. *Applied smoothing techniques for data analysis*. New York: Oxford University Press Inc; 1997.
- [40] Hill PD. Kernel estimation of a distribution function. *Communications in Statistics - Theory and Methods*. 1985;14, 3:605–620.
- [41] Jones MC. Simple boundary correction for kernel density estimation. *Statistics and Computing*. 1993;3, 3:135–146.
- [42] Silverman BW. *Density estimation for statistics and data analysis*. London: Chapman & Hall; 1986.
- [43] Cassini JD. De l'origine et du progrès de l'astronomie et de son usage dans la géographie et dans la navigation. L'Imprimerie Royale. 1693;:36.
- [44] Bos HJM. *The lemniscate of bernoulli*. Dordrecht: Springer Netherlands; 1974.