



Brief paper

Information value in nonparametric Dirichlet-process Gaussian-process (DPGP) mixture models[☆]



Hongchuan Wei^a, Wenjie Lu^b, Pingping Zhu^c, Silvia Ferrari^{d,1}, Miao Liu^e,
Robert H. Klein^f, Shayegan Omidshafiei^g, Jonathan P. How^{h,2}

^a Duke University, 104 misty woods circle, apt L, 27514 chapel hill, United States

^b Duke University, CB11.09.203, 15 Broadway, Ultimo NSW 2007, Australia

^c Cornell University, 529 Upson Hall, 14853, Ithaca, United States

^d Cornell University, 543 Upson Hall, 14853, Ithaca, United States

^e Massachusetts Institute of Technology, 02139 Cambridge, United States

^f Massachusetts Institute of Technology, 1970 Harvard Ave E, Apt 1, Seattle, WA 98102, United States

^g Massachusetts Institute of Technology, 02139 Cambridge, United States

^h Massachusetts Institute of Technology, Building 33, Room 326, 02139 Cambridge, United States

ARTICLE INFO

Article history:

Received 8 July 2015

Received in revised form

27 February 2016

Accepted 22 June 2016

Keywords:

Information theory

Bayesian nonparametric models

Gaussian process

Dirichlet process

Information gain

ABSTRACT

This paper presents tractable information value functions for Dirichlet-process Gaussian-process (DPGP) mixture models obtained via collocation methods and Monte Carlo integration. Quantifying information value in tractable closed form is key to solving control and estimation problems for autonomous information-gathering systems. The properties of the proposed value functions are analyzed and then demonstrated by planning sensor measurements so as to minimize the uncertainty in DPGP target models that are learned incrementally over time. Simulation results show that sensor planning based on expected KL divergence outperforms algorithms based on mutual information, particle filters, and randomized methods.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

Bayesian nonparametric models, such as the Dirichlet-process Gaussian-process (DPGP) mixtures, have been recently developed for modeling multiple dynamic processes adaptively from data. When data becomes available over time, DPGP clusters and parameters are expanded or compacted incrementally, as needed, to avoid growing the model dimensionality indefinitely as the size

of the database increases (Rasmussen, 1999). Because of these characteristics, DPGP mixtures are particularly useful in life-long learning and sensing problems, and present the opportunity for planning the measurement sequence so as to optimize the value of future data.

To date, DPGPs have been shown effective at modeling traffic patterns (Joseph, Doshi-Velez, Huang, & Roy, 2011), clinical identification (Ross & Dy, 2013), and gene expression time series analysis (Hensman, Rattray, & Lawrence, 2015). Although information theoretic functions have become a common approach for representing information value in sensing and control problems (Lu, Zhang, & Ferrari, 2014), they are not directly applicable to DPGPs because they are defined in terms of finite-dimensional probability distributions (Rasmussen, 1999). Previous approaches to quantifying information value in GPs consist of inferring the process state and then determining the differential entropy or mutual information between the state and the measurements (Krause, Singh, & Guestrin, 2008). Similar work in Ny and Pappas (2009) considered the mutual information between the state of a moving target and available measurements using a GP model. Previous work by the

[☆] This work was supported by ONR MURI Grant N000141110688. The material in this paper was partially presented at the International Conference on Intelligent Robots and Systems (IROS), September 14–18, 2014, Chicago, IL, USA. This paper was recommended for publication in revised form by Associate Editor Andrey V. Savkin under the direction of Editor Ian R. Petersen.

E-mail addresses: hongchuan.wei@duke.edu (H. Wei), wenjie.lu@duke.edu (W. Lu), pingping.zhu@cornell.edu (P. Zhu), ferrari@cornell.edu (S. Ferrari), miaoliu@mit.edu (M. Liu), rhklein@mit.edu, bobklein2@gmail.com (R.H. Klein), shayegan@mit.edu (S. Omidshafiei), jhow@mit.edu (J.P. How).

¹ Fax: +1 607255 1222.

² Fax: +1 617 253 7397.

authors proposed using the Kullback–Leibler (KL) divergence obtained by approximating the GPs by probability distributions about a set of collocation points for a given DP prior (Wei, Lu, & Ferrari, 2012; Wei et al., 2014). However, these GP-based methods are not applicable to the DPGP mixture models considered in this paper.

Building on the collocation point method in Wei et al. (2012), this paper presents a systematic approach for obtaining tractable information value functions, such as the expected KL divergence, for DPGP models of multiple dynamic targets. New conditional and expected information theoretic functions for DPGPs are derived in Sections 4.2–4.3 to quantify information value in non-myopic sequential observations. New theoretical results (Section 4.4) are used to obtain a computationally efficient approximation of DPGP information value via Monte Carlo integration. The new analysis in Section 4.4 proves that this approximation is an unbiased estimator of the original DPGP value function, and is characterized by an error covariance that decreases linearly with the number of samples. As a result, it can be used to control moving sensors in real time while observing dynamic targets. Simulation results show that using this approximate DPGP KL divergence for sensor planning is far more effective than using DPGP mutual information, particle filtering (Wang, Yin, & Hong, 2008), or randomized methods (Fulgenzi, Tay, Spalanzani, & Laugier, 2008).

2. Problem formulation and assumptions

DPGP mixtures have been recently shown to be a new and useful paradigm for modeling multiple dynamic processes from data, as exemplified by traffic patterns modeled from measurements obtained by a helicopter flying over the greater Boston area (Joseph et al., 2011). The purpose of the information functions derived and analyzed in this paper is to provide design objectives for controlling moving sensors such that their measurement value is optimized over time. Consider the problem of modeling N unknown and independent targets in a workspace $\mathcal{W} \subset \mathbb{R}^2$, using multiple sensor measurements obtained over time. Assume each target can be described by a time-invariant nonlinear ordinary differential equation (ODE),

$$\dot{\mathbf{x}}_j(t) = \mathbf{f}_i[\mathbf{x}_j(t)] \triangleq \mathbf{v}_j(t), \quad j = 1, \dots, N(t) \quad (1)$$

where the vector function $\mathbf{f}_i : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, referred to as a *velocity field* (VF), is drawn from a set of continuously differentiable velocity fields, $\mathcal{F} = \{\mathbf{f}_1, \dots, \mathbf{f}_M\}$, $\mathbf{x}_j \in \mathcal{W}$ denotes the j th target position, $\mathbf{v}_j \in \mathbb{R}^2$ denotes the j th target velocity, and $t \in [t_0, t_f]$. The number of targets N in \mathcal{W} is unknown, time varying, and may not equal M .

At the initial time, t_0 , no prior target information is available and, thus, M and \mathcal{F} are to be learned from data. The data consists of position and velocity measurements,

$$\mathbf{m}_j(t) \triangleq \begin{bmatrix} \mathbf{y}_j(t) \\ \mathbf{z}_j(t) \end{bmatrix} = \begin{bmatrix} \mathbf{x}_j(t) \\ \mathbf{v}_j(t) \end{bmatrix} + \mathbf{n}(t), \quad j = 1, \dots, N(t) \quad (2)$$

for $\mathbf{x}_j(t) \in \mathcal{S}(t)$, $t > t_0$, where $\mathcal{S}(t) \subset \mathcal{W}$ is the moving sensor field-of-view (FoV), and $\mathbf{n} \in \mathbb{R}^4$ is an additive Gaussian measurement error with zero mean and known standard deviations σ_x and σ_v (Wei, Ross, Varisco, Krief, & Ferrari, 2013). If $\mathbf{x}_j(t) \notin \mathcal{S}(t)$, $\mathbf{m}_j(t)$ belongs to an empty set. The measurement–target association is carried out by established data association algorithms (Bar-Shalom & Blair, 2000; Bar-Shalom & Tse, 1975; Pattipati, Deb, Bar-Shalom, & Washburn, 1992), and any errors are assimilated in the DPGP model uncertainty representation.

From (1), it can be seen that a VF, \mathbf{f}_i , projects a target position, \mathbf{x}_j , to a velocity, \mathbf{v}_j , and, thus, can be viewed as a two-dimensional spatial phenomenon. Because \mathbf{f}_i is unknown and is not necessarily drawn uniformly at random from \mathcal{F} , the association of a target

with a VF in \mathcal{F} can be represented by a discrete random variable, G_j , with a range $\mathcal{I} = \{1, \dots, M\}$ that is possibly infinite. Then, the event $\{G_j = i\}$ represents the association of target j with the velocity field $\mathbf{f}_i \in \mathcal{F}$, as shown in (1), and the N targets can be modeled as a multioutput GP via nonlinear regression, as follows.

A multioutput GP defines a multivariate distribution over functions, $P(\mathbf{f}_i)$, where $\mathbf{f}_i : \mathcal{W} \rightarrow \mathbb{R}^n$, and here $n = 2$ and $i = 1, \dots, M$ (Joseph et al., 2011). Let $F = \{\mathbf{f}_i(\mathbf{x}_1), \dots, \mathbf{f}_i(\mathbf{x}_N) \mid \mathbf{x}_N \in \mathcal{W}\}$ be a set of vector function values evaluated at N points in \mathcal{W} . Then, $P(\mathbf{f}_i)$ is a multioutput GP if for any finite subset $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ the marginal distribution $P(F)$ is a joint multivariate Gaussian distribution (Rasmussen & Williams, 2006). A multioutput GP is specified by its mean vector function,

$$\theta_i(\mathbf{x}_j) = \mathbb{E}_{\mathbf{v}_j}[\mathbf{f}_i(\mathbf{x}_j)], \quad \forall \mathbf{x}_j \in \mathcal{W} \quad (3)$$

and its covariance matrix function,

$$\Psi_i(\mathbf{x}_j, \mathbf{x}'_j) = \mathbb{E}_{\mathbf{v}_j} \{ [\mathbf{f}_i(\mathbf{x}_j) - \theta_i(\mathbf{x}_j)][\mathbf{f}_i(\mathbf{x}'_j) - \theta_i(\mathbf{x}'_j)]^T \} \quad (4)$$

for any $\mathbf{x}_j, \mathbf{x}'_j \in \mathcal{W}$, where $\mathbb{E}_{\mathbf{v}_j}[\cdot]$ denotes the expectation operator with respect to \mathbf{v}_j (Rasmussen & Williams, 2006). Assuming stationary covariance functions, the notation,

$$\mathbf{f}_i(\mathbf{x}_j) \sim \text{GP}_i(\theta_i(\mathbf{x}_j), \Psi_i(\mathbf{x}_j, \mathbf{x}'_j)), \quad \forall \mathbf{x}_j, \mathbf{x}'_j \in \mathcal{W} \quad (5)$$

or in short $\mathbf{f}_i(\mathbf{x}_j) \sim \text{GP}_i$, can be used to indicate that \mathbf{f}_i is “distributed as” the Gaussian process GP_i , for $i = 1, \dots, M$. For simplicity, in this paper it is assumed that the elements of \mathbf{v}_j are independent, such that Ψ_i is a diagonal and positive-definite matrix. Also, it is assumed that the M Gaussian processes, $\text{GP}_1, \dots, \text{GP}_M$, share the same covariance matrix $\Psi_i = \Psi$ and, therefore, $\text{GP}_i = \text{GP}(\theta_i, \Psi)$, where Ψ is known *a priori*.

For N independent targets, the target–VF association variables, G_1, \dots, G_N , can be assumed to be independent and identically distributed (i.i.d.), such that the probability of event $\{G_j = i\}$ is

$$P(G_j = i) = \pi_i, \quad \forall j. \quad (6)$$

Letting $\boldsymbol{\pi} \triangleq [\pi_1 \dots \pi_M]^T$ denote a prior probability vector that satisfies the properties, $\sum_{i=1}^M \pi_i = 1$, and $\pi_i \in [0, 1], \forall i \in \mathcal{I}$, the prior distribution of G_j can be represented by an M -dimensional categorical distribution $\text{Cat}(\boldsymbol{\pi})$, with probability mass function (PMF) $\boldsymbol{\pi}$ (Bogachev & Ruas, 2007). Then, the tuple $\{\mathcal{F}, \boldsymbol{\pi}\}$ provides sufficient statistics for modeling all N targets from data.

3. DPGP target modeling and planning

When multiple targets obey the same VF, a prior distribution on $\boldsymbol{\pi}$ is required to provide the ability of learning the clustering of target behaviors from sensor measurements. Dirichlet processes (DPs) have been successfully applied to the clustering of data into an unknown number of clusters because they allow for the creation and deletion of clusters, as necessary, while new data is obtained over time. Let (A, B) be a measurable space, where B is a σ -algebra on a set A (Bogachev & Ruas, 2007). A finite measurable partition $\{B_i\}_{i=1}^n$ of A is a collection of sets $B_i \in B$, such that $B_i \cap B_j = \emptyset$, if $i \neq j$; and $\cup_{i=1}^n B_i = A$. Let H be a finite non-zero measure on the measurable space (A, B) , and let α be a positive real number. A DP with parameters H and α , denoted by $\text{DP}(\alpha, H(A))$, is the distribution of a random probability measure P if, for any finite measurable partition $\{B_i\}_{i=1}^n$ of A , the following holds,

$$[P(B_1) \dots P(B_n)]^T \sim \text{Dir}(\alpha H(B_1), \dots, \alpha H(B_n)) \quad (7)$$

where “Dir” denotes the Dirichlet distribution (Teh, Jordan, Beal, & Blei, 2006).

From Section 2, the base distribution H is a GP with zero mean, or $\text{GP}_0 = \text{GP}(\mathbf{0}, \Psi)$, and \mathcal{A} is the space of continuously differentiable functions $C^1(\mathcal{W})$. Then, the following DPGP mixture (Joseph et al., 2011),

$$\begin{aligned} \{\theta_i, \pi\} &\sim \text{DP}(\alpha, \text{GP}_0), \quad i = 1, \dots, \infty \\ G_j &\sim \text{Cat}(\pi), \quad j = 1, \dots, N \\ \mathbf{f}_{G_j}(\mathbf{x}) &\sim \text{GP}(\theta_{G_j}, \Psi), \quad \mathbf{x} \in \mathcal{W}, \quad j = 1, \dots, N, \end{aligned} \quad (8)$$

can be used to model the dynamic targets (1) from sensor measurements (2) that are obtained incrementally over a period of time $[t_0, t_f]$.

When new sensor measurements become available, (8) can be updated to make use of the latest target information for tracking, pursuit, or estimation. Let Δt denote a small but finite constant time interval required by the sensor to obtain and process new target measurements using a finite sampling rate. Then, $\mathbf{m}_j(k)$ denotes sensor measurements obtained from the j th target during the k th sampling interval, $[t_k, t_k + \Delta t)$. Since the sensor FoV is stationary during sampling, $\mathbf{m}_j(k)$ is known provided $\mathbf{x}_j(t) \in \mathcal{S}(k)$ for some $t \in [t_k, t_k + \Delta t)$. Because DPGP cluster learning depends on all targets observed over time, all measurements obtained up to time k , denoted by $M(k) \triangleq M_1(k) \cup \dots \cup M_N(k)$, where $M_j(k) \triangleq \{\mathbf{m}_j(\ell) \mid 1 \leq \ell \leq k\}$, are used for the DPGP update at time k . Assume the DP prior is fixed during Δt . Then, the VF clusters can remain unchanged and the VF-target associations learned at time k , $\mathcal{G}(k) \triangleq \{G_j(k) \mid 1 \leq j \leq N\}$, are also part of the measurement database, $Q(k) \triangleq \{M(k), \mathcal{G}(k)\}$, used to update (8) at time k . Furthermore, from $Q(k)$, new sensor measurements can be planned at t_{k+1}, t_{k+2}, \dots , so as to minimize the DPGP model uncertainty in future updates.

4. Information value in DPGP mixture models

This section presents a novel approach for evaluating information value in DPGP mixture models so that the reduction in uncertainty brought about by new sensor measurements can be estimated from past data and a prior DPGP model. Evaluating information value in closed form is important both for assessing the benefit of a given database and for planning sensor motions and parameters. In order to derive information functions for the DPGP in (8), the M VFs in (1) are evaluated at a finite number of collocation points in the domain of integration \mathcal{W} of the ODEs (1). The approach is illustrated for the KL divergence because it is typically most effective for solving sensor planning problems (Kastella, 1997; Wei et al., 2012; Zhang, Ferrari, & Cai, 2012), but can easily be extended to other information theoretic functions, such as mutual information (Section 5).

4.1. DPGP KL divergence

Let $p(x)$ and $q(x)$ denote two known probability density functions (PDFs) of a continuous random variable $x \in \mathbb{R}$. Then, the KL divergence or relative entropy,

$$D(p(x) \parallel q(x)) = \int_{-\infty}^{\infty} p(x) \ln \frac{p(x)}{q(x)} dx \quad (9)$$

can be used to represent the “distance” between $p(x)$ and $q(x)$. Although it does not constitute a true distance metric because it is nonadditive, nonsymmetric, and does not obey the triangle inequality, the KL divergence can represent the change in a PDF, e.g., due to a new measurement or observation (Cover & Thomas, 1991).

Let $\xi_l \in \mathcal{W}$ denote the l th collocation point chosen from a uniform grid of L points in \mathcal{W} grouped in a $2L \times 1$ vector $\xi = [\xi_1^T \dots \xi_L^T]^T$, as in basic collocation methods (Nobile, Tempone, & Webster, 2008). Each ODE in (1) is discretized about every

collocation point by evaluating \mathbf{f}_i at ξ_l for $l = 1, \dots, L$, such that,

$$\mathbf{v}_i \triangleq [\mathbf{f}_i(\xi_1)^T \dots \mathbf{f}_i(\xi_L)^T]^T = \mathbf{v}_i(\xi) \quad (10)$$

approximates \mathbf{f}_i in \mathcal{W} . Similarly, all M velocity fields in \mathcal{F} are discretized and grouped into a $2LM \times 1$ vector of random velocity variables $\mathbf{v} \triangleq [\mathbf{v}_1(\xi)^T \dots \mathbf{v}_M(\xi)^T]^T$.

Now, let $p(\mathbf{v})$ and $q(\mathbf{v})$ denote two joint PDFs of the $2LM$ elements of the random velocity vector \mathbf{v} , where each PDF is obtained from the DPGP model (8). Then, the “distance” between the two DPGP parameterizations can be represented by the DPGP KL divergence

$$D(p(\mathbf{v}) \parallel q(\mathbf{v})) \triangleq \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} p(\mathbf{v}) \ln \frac{p(\mathbf{v})}{q(\mathbf{v})} d\mathbf{v}. \quad (11)$$

In order to evaluate (11), the joint PDF of $p(\mathbf{v})$ needs to be expressed in terms of the DPGP parameters in (8). Given $\{\pi, \theta_i, \Phi\}$, the random vector \mathbf{v}_i has a mixed multivariate Gaussian distribution with mean and covariance matrix calculated from θ_i and Φ . However, the number of components in the mixture model is infinite. Therefore, computing (11) is computationally very expensive due to the multiple integrals of the joint PDFs. Thus, the next subsections present several steps by which the KL divergence can be simplified and approximated by an information function that is computationally efficient, using conditional independence assumptions and Monte Carlo integration.

4.2. Conditional DPGP KL divergence

For a non-myopic process, the information value is to be conditioned on all prior measurements or observations. Thus, as a first step, the KL divergence is obtained in terms of conditional probability densities that take into account the measurement database $Q(k)$ available or given at time t_k . The reduction in DPGP model uncertainty brought about by a measurement $\mathbf{m}_j(k+1)$ can be represented by the conditional KL divergence,

$$\begin{aligned} D(\mathbf{v}; \mathbf{m}_j(k+1) | Q(k)) \\ \triangleq D(p(\mathbf{v} | \mathbf{m}_j(k+1), Q(k)) \parallel p(\mathbf{v} | Q(k))) \end{aligned} \quad (12)$$

in terms of the joint PDFs of \mathbf{v} obtained from (8). Because the VF vectors of two targets are conditionally independent given $Q(k)$, the joint PDF is factorized as,

$$p(\mathbf{v} | Q(k)) = p(\mathbf{v}_1, \dots, \mathbf{v}_M | Q(k)) = \prod_{i=1}^M p(\mathbf{v}_i | Q(k)) \quad (13)$$

where each joint PDF $p(\mathbf{v}_i | Q(k))$ can be obtained from the GP covariance and mean as follows.

From GP regression (Rasmussen & Williams, 2006), given a database $Q(k)$, \mathbf{v}_i is characterized by the multivariate joint Gaussian PDF,

$$p(\mathbf{v}_i | Q(k)) = \frac{-\frac{1}{2}[\mathbf{v}_i - \boldsymbol{\mu}_i(k)]^T [\boldsymbol{\Sigma}_i(k)]^{-1} [\mathbf{v}_i - \boldsymbol{\mu}_i(k)]}{\sqrt{(2\pi)^L |\boldsymbol{\Sigma}_i(k)|}}$$

where $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$ are the mean vector and covariance matrix of the measurements associated with the i th VF, respectively, and $|\cdot|$ denotes the matrix determinant. Both $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$ can be computed from the cross-covariance matrix of the GP regression,

$$\begin{aligned} \Phi(\mathbf{X}, \mathbf{X}') &\triangleq \mathbb{E}_{\mathbf{v}_i, \mathbf{v}'_i} [(\mathbf{v}_i - \mathbb{E}[\mathbf{v}_i])(\mathbf{v}'_i - \mathbb{E}[\mathbf{v}'_i])^T] \\ &= \begin{bmatrix} \Psi(\mathbf{x}_1, \mathbf{x}'_1) & \dots & \Psi(\mathbf{x}_1, \mathbf{x}'_n) \\ \vdots & \ddots & \vdots \\ \Psi(\mathbf{x}_m, \mathbf{x}'_1) & \dots & \Psi(\mathbf{x}_m, \mathbf{x}'_n) \end{bmatrix} \end{aligned} \quad (14)$$

where $\mathbf{X} = [\mathbf{x}_1^T \dots \mathbf{x}_m^T]^T$ and $\mathbf{X}' = [(\mathbf{x}'_1)^T \dots (\mathbf{x}'_n)^T]^T$ for any $\mathbf{x}_j, \mathbf{x}'_l \in \mathbb{R}^2$. From (10), $\mathbf{v}_i \triangleq [\mathbf{f}_i(\mathbf{x}_1)^T \dots \mathbf{f}_i(\mathbf{x}_m)^T]^T$ and $\mathbf{v}'_i \triangleq [\mathbf{f}_i(\mathbf{x}'_1)^T$

$\dots \mathbf{f}_i(\mathbf{x}'_n)^T]^T$ are obtained by evaluating \mathbf{f}_i at every point in \mathbf{X} and \mathbf{X}' , respectively.

Now, let $\mathbf{Y}_i(k) = [\mathbf{y}_1^T(\cdot) \ \mathbf{y}_2^T(\cdot) \ \dots]^T$ and $\mathbf{Z}_i(k) = [\mathbf{z}_1^T(\cdot) \ \mathbf{z}_2^T(\cdot) \ \dots]^T$ denote two vectors containing all position and velocity measurements associated with \mathbf{f}_i up to time k , respectively. Measurements $\mathbf{y}_j(l)$, $\mathbf{z}_j(l)$ are associated with a velocity field $\mathbf{f}_i \in \mathcal{F}$ by the index set $\mathcal{G}_j(k)$ learned using (8). Let $\mathbf{A} \triangleq \Phi[\mathbf{Y}_i(k), \mathbf{Y}_i(k)] + \sigma_v^2 \mathbf{I}_{2k}$ denote the cross-covariance matrix obtained by GP regression (Rasmussen & Williams, 2006). Then, the measurement mean vector and covariance matrix for a collocation point vector ξ are

$$\boldsymbol{\mu}_i(k) = \Phi[\xi, \mathbf{Y}_i(k)] \mathbf{A}^{-1} \mathbf{Z}_i(k) \triangleq \boldsymbol{\mu}_{i,k} \quad (15)$$

$$\boldsymbol{\Sigma}_i(k) = \Phi(\xi, \xi) - \Phi[\xi, \mathbf{Y}_i(k)] \mathbf{A}^{-1} \Phi[\mathbf{Y}_i(k), \xi] \triangleq \boldsymbol{\Sigma}_{i,k}.$$

Without loss of generality, consider $G_j(k+1) = i$. Then, since measurements obtained from one VF do not contain information about other VFs, it follows that,

$$D(\mathbf{v}; \mathbf{m}_j(k+1)|Q(k)) = D(\mathbf{v}_i; \mathbf{m}_j(k+1)|Q(k)) \quad (16)$$

as proven in the Appendix. Now, if $\mathbf{m}_j(k+1)$ and $G_j(k+1)$ are both given, \mathbf{v}_i is characterized by the multivariate joint Gaussian PDF in (13), and the conditional DPGP KL divergence (12) reduces to the closed form,

$$D(\mathbf{v}; \mathbf{m}_j(k+1)) = \frac{1}{2} \left[\text{tr}(\boldsymbol{\Sigma}_{i,k}^{-1} \boldsymbol{\Sigma}_{i,k+1}) - \ln \left(\frac{|\boldsymbol{\Sigma}_{i,k+1}|}{|\boldsymbol{\Sigma}_{i,k}|} \right) \right] + \frac{1}{2} (\boldsymbol{\mu}_{i,k+1} - \boldsymbol{\mu}_{i,k})^T \boldsymbol{\Sigma}_{i,k}^{-1} (\boldsymbol{\mu}_{i,k+1} - \boldsymbol{\mu}_{i,k}) - L \quad (17)$$

where $\text{tr}(\cdot)$ denotes the trace of a matrix, and, for brevity, conditioning on $Q(k)$ is omitted here and in the remainder of the paper. When $\mathbf{m}_j(k+1)$ is unknown, as in sensor planning problems, the expected DPGP KL divergence can be derived as shown in the next subsection.

4.3. Expected DPGP KL divergence

When planning the motion of a sensor, such as a camera mounted on a helicopter (Joseph et al., 2011), the value of $\mathbf{m}_j(k+1)$ is unknown, and $G_j(k+1)$ cannot be learned from data. In this case, (17) cannot be evaluated because at time t_{k+1} the measurement mean and covariance in (15) are unknown. Then, the DPGP information value can be estimated by taking the expectation of (12) with respect to $\mathbf{m}_j(k+1)$ and $G_j(k+1)$ as follows

$$\hat{D}(\mathbf{v}; \mathbf{m}_j(k+1)) \triangleq \mathbb{E}_{\mathbf{m}_j} \{ \mathbb{E}_{G_j} \{ D(\mathbf{v}; \mathbf{m}_j(k+1)) \} \}. \quad (18)$$

Let L_{ij} denote the likelihood of event $\{G_j(k+1) = i\}$, such that, $L_{ij} \triangleq p(M_j(k)|M_j^c(k), G_j(k+1) = i)$, where $M_j^c = Q \setminus M_j$ denotes the complement set of M_j in Q . Then, the posterior probability of $\{G_j(k+1) = i\}$ is,

$$w_{ij} \triangleq p(G_j(k+1) = i|Q(k)) = \frac{\pi_i L_{ij}}{\sum_{i=1}^M \pi_i L_{ij}} \quad (19)$$

where $\pi_i = p(G_j(k+1) = i)$, defined in (6), is the prior probability that the j th target obeys \mathbf{f}_i .

From (16) and (18), the DPGP expected KL divergence (EKLD) can be written as,

$$\hat{D}(\mathbf{v}; \mathbf{m}_j(k+1)) = \sum_{i=1}^M w_{ij} \hat{D}(\mathbf{v}_i; \mathbf{m}_j(k+1)) \quad (20)$$

and evaluated using the likelihood,

$$L_{ij} = \prod_{\ell=1}^k \frac{\exp \left\{ -\frac{1}{2} [\mathbf{z}_{j,\ell} - \boldsymbol{\eta}_{j,\ell}]^T [\boldsymbol{\Lambda}_{j,\ell}]^{-1} [\mathbf{z}_{j,\ell} - \boldsymbol{\eta}_{j,\ell}] \right\}}{2\pi |\boldsymbol{\Lambda}_{j,\ell}|}$$

obtained from the i th GP component of the latest DPGP model, where $\mathbf{z}_{j,\ell} = \mathbf{z}_j(\ell)$, and the target velocity mean, $\boldsymbol{\eta}_{j,\ell} = \boldsymbol{\eta}_j(\ell)$, and covariance, $\boldsymbol{\Lambda}_{j,\ell} = \boldsymbol{\Lambda}_j(\ell)$, at $\mathbf{y}_j(\ell)$ are calculated by replacing ξ with $\mathbf{y}_j(\ell)$ in (15).

Then, the expected KL divergence of \mathbf{v}_i in (20) is obtained by marginalizing the original KL divergence function over all possible values of $\mathbf{m}_j(k+1)$,

$$p(\mathbf{m}_j(k+1)|G_j(k+1) = i) = p(\mathbf{z}_j(k+1)|\mathbf{x}_j(k+1)) p(\mathbf{x}_j(k+1)|G_j(k+1) = i) \quad (21)$$

where $p(\mathbf{z}_j(k+1)|\mathbf{x}_j(k+1))$ can be calculated from (2). Using Euler integration and the ODE (1), it follows that,

$$p(\mathbf{x}_j(k+1)|G_j(k+1) = i) = \int_{\mathbb{R}^2} p(\mathbf{v}_j(k)|\mathbf{x}_j(k)) f_X(\mathbf{x}_j(k+1) - \mathbf{v}_j(k)\Delta t) d\mathbf{v}_j(k) \quad (22)$$

where $f_X(\cdot)$ is the PDF of $\mathbf{x}_j(k)$ obtained via filtering techniques (Wei et al., 2014). The conditional joint probability distribution of the target speed, $p(\mathbf{v}_j(k)|\mathbf{x}_j(k))$, is a multivariate Gaussian distribution with mean and covariance calculated from (15) by replacing ξ with $\mathbf{x}_j(k)$. Thus, the information value of a future measurement $\mathbf{m}_j(k+1)$ can be estimated using Eqs. (19)–(22).

When multiple target measurements can be obtained during one sampling interval $[t_k, t_k + \Delta t)$, say $\mathbf{m}(k) = [\mathbf{m}_j^T(k) \ \mathbf{m}_i^T(k) \ \dots]^T$, the DPGP-EKLD is given by,

$$\begin{aligned} \hat{D}(\mathbf{v}; \mathbf{m}(k+1)) &= \sum_{j, \mathbf{x}_j \in \mathcal{S}(k+1)} \hat{D}(\mathbf{v}; \mathbf{m}_j(k+1)) \\ &= \sum_{j,i} w_{ij} \int_{\mathcal{S}(k+1)} \int_{\mathbb{R}^2} D(\mathbf{v}_i; \mathbf{m}_j(k+1)) p(\mathbf{z}_j(k+1)|\mathbf{x}_j(k+1)) \\ &\quad \times p(\mathbf{x}_j(k+1)|G_j(k+1) = i) d\mathbf{z}_j(k+1) d\mathbf{x}_j(k+1) \end{aligned} \quad (23)$$

where all quantities are defined and calculated as shown in (19)–(22). Because the DPGP-EKLD function in (23) involves a 6th-order integral, its computation is typically very burdensome. Using the methodology presented in the next subsection, it is possible to simplify the DPGP-EKLD to a double integral, a form that can be implemented in real-time for sensor planning (Section 5).

4.4. DPGP expected KL divergence approximation

The theoretical results presented in this subsection allow for the analytical approximation of the inner double integral in the DPGP-EKLD function (23), and provide performance guarantees for the subsequent Monte Carlo (MC) integration required to evaluate the DPGP-EKLD approximation. Since the probability $p(\mathbf{x}_j(k+1)|G_j(k+1) = i)$ is constant with respect to $\mathbf{z}_j(k+1)$, (23) can be simplified to,

$$\begin{aligned} \hat{D}(\mathbf{v}; \mathbf{m}(k+1)) &= \sum_{j,i} w_{ij} \int_{\mathcal{S}(k+1)} h_i[\mathbf{x}_j(k+1)] \\ &\quad \times p(\mathbf{x}_j(k+1)|G_j(k+1) = i) d\mathbf{x}_j(k+1) \end{aligned} \quad (24)$$

where

$$h_i[\mathbf{x}_j(k+1)] \triangleq \int_{\mathbb{R}^2} \left\{ D(\mathbf{v}_i; \mathbf{m}_j(k+1)) \times p(\mathbf{z}_j(k+1)|\mathbf{x}_j(k+1)) \right\} d\mathbf{z}_j(k+1). \quad (25)$$

Then, the above second-order integral can be solved analytically according to the following proposition.

Proposition 1. Consider a Gaussian process GP_i with known covariance matrix function Ψ_i , and assume the target position is $\mathbf{x}_j(k+1)$. Then, the second-order integral in (25) affords the analytical solution,

$$h_i[\mathbf{x}_j(k+1)] = \frac{1}{2} \left[\text{tr}(\Sigma_{i,k}^{-1} \Sigma_{i,k+1}) - \ln(|\Sigma_{i,k+1} \Sigma_{i,k}^{-1}|) - 2L + \text{tr}(\mathbf{Q}^{-1} \mathbf{R}^T \Sigma_{i,k}^{-1} \mathbf{R} \mathbf{Q}^{-1}) \sigma_v^2 \right] \quad (26)$$

where $\Sigma_{i,k}$ and $\Sigma_{i,k+1}$ are the velocity covariance matrices at times k and $k+1$, respectively, and,

$$\mathbf{A} \triangleq \Phi[\mathbf{Y}_i(k), \mathbf{Y}_i(k)] + \sigma_v^2 \mathbf{I}_{2k} \quad (27)$$

$$\mathbf{B} \triangleq \Phi[\mathbf{Y}_i(k), \mathbf{x}_j(k+1)] \quad (28)$$

$$\mathbf{D} \triangleq \Phi[\mathbf{x}_j(k+1), \mathbf{x}_j(k+1)] + \sigma_v^2 \mathbf{I}_2 \quad (29)$$

$$\mathbf{R} \triangleq \Phi[\xi, \mathbf{x}_j(k+1)] - \Phi[\xi, \mathbf{Y}_i(k)] \mathbf{A}^{-1} \mathbf{B} \quad (30)$$

$$\mathbf{Q} \triangleq \mathbf{D} - \mathbf{B}^T \mathbf{A}^{-1} \mathbf{B} \quad (31)$$

for any $\mathbf{x}_j(k+1) \in \mathcal{S}(k+1)$, where ξ is a vector of collocation points, \mathbf{Y}_i is a vector of all past position measurements, and Φ is the cross-covariance matrix of GP_i .

Proof. Using the matrix inversion lemma, $[\Sigma_{i,k+1}]^{-1}$ can be written as,

$$\begin{bmatrix} \mathbf{A}^{-1}(\mathbf{I}_{2k} + \mathbf{B} \mathbf{Q}^{-1} \mathbf{B}^T \mathbf{A}^{-1}) & -\mathbf{A}^{-1} \mathbf{B} \mathbf{Q}^{-1} \\ \mathbf{Q}^{-1} \mathbf{B}^T \mathbf{A}^{-1} & \mathbf{Q}^{-1} \end{bmatrix}$$

where \mathbf{A} , \mathbf{B} , \mathbf{D} , and \mathbf{Q} are defined from (27)–(31). Substituting $\Sigma_{i,k+1}^{-1}$ in the GP mean vector expression (15), the change in the mean during two consecutive time steps can be written as,

$$\boldsymbol{\mu}_i(k+1) - \boldsymbol{\mu}_i(k) = \mathbf{R} \mathbf{Q}^{-1} \mathbf{z}'_j(k+1) \quad (32)$$

where $\mathbf{z}'_j(k+1) \triangleq \mathbf{z}_j(k+1) - \mathbb{E}[\mathbf{z}_j(k+1)]$ is a Gaussian random vector with zero mean and covariance matrix $\sigma_v^2 \mathbf{I}_2$. Substituting (32) into (25), the integral (26) can be simplified to

$$h_i[\mathbf{x}_j(k+1)] = \frac{1}{2} \left[\text{tr}(\Sigma_{i,k}^{-1} \Sigma_{i,k+1}) - \ln(|\Sigma_{i,k+1} \Sigma_{i,k}^{-1}|) - 2L + \text{tr}(\mathbf{Q}^{-1} \mathbf{R}^T \Sigma_{i,k}^{-1} \mathbf{R} \mathbf{Q}^{-1}) \sigma_v^2 \right]. \quad \square$$

Even with the above simplification, the computational complexity of evaluating $h_i[\mathbf{x}_j(k+1)]$ is $O(L^3 + k^3)$, where L is the number of collocation points and k is the time index. Thus, it may be infeasible to evaluate (24) for all possible measurements. Because the integrand of (24) goes to zero when $p(\mathbf{x}_j(k+1)|G_j(k+1)=i)$ goes to zero, and $h_i[\mathbf{x}_j(k+1)]$ is finite, this computation can be significantly reduced by the approach known as Monte Carlo integration (Press, 2007). Let $\xi^{(1)}, \dots, \xi^{(S)}$ denote S values of $\mathbf{x}_j(k+1)$ drawn identically and independently from $p(\mathbf{x}_j(k+1)|G_j(k+1)=i)$ in (22). Then, the integral in (24) can be computed numerically by evaluating its integrand at each sample with nonzero probability, i.e.,

$$\hat{D}(v; \mathbf{m}(k+1)) \approx \sum_{j,i} \frac{w_{ij}}{S} \sum_{l=1}^S h_i(\xi^{(l)}) \mathbf{1}_{\mathcal{S}(k+1)}(\xi^{(l)}) \quad (33)$$

where,

$$\mathbf{1}_{\mathcal{S}(k+1)}(\xi^{(l)}) \triangleq \begin{cases} 1, & \xi^{(l)} \in \mathcal{S}(k+1) \\ 0, & \xi^{(l)} \notin \mathcal{S}(k+1) \end{cases} \quad (34)$$

is the indicator function. Similarly to the collocation points used to discretize the DPGP information value, these samples represent points in \mathcal{W} . But, unlike collocation points, which are chosen from

a uniform grid, the Monte Carlo integration samples are chosen by sampling a known distribution.

The remainder of this subsection proves that, based on the following three lemmas, the approximation in (33) is an unbiased estimator of the DPGP-EKLD in (24), and its error variance decreases linearly with S .

Lemma 2. The velocity covariance matrix \mathbf{Q} , defined in (31), obeys the element-wise inequality $\Psi_0 + \{\sigma_v^2/[k \text{tr}(\Psi_0) + \sigma_v^2]\} \mathbf{I}_2 \leq \mathbf{Q} \leq \Psi_0 + \sigma_v^2 \mathbf{I}_2$, where $\Psi_0 = \Psi(\mathbf{0}, \mathbf{0})$ is a constant matrix obtained by evaluating the stationary covariance matrix function at zero, $\sigma_v \in \mathbb{R}$ is the standard deviation of the measurement noise of velocity, and $k \in \mathbb{Z}_+$ is the time index.

Proof. From (27), the matrix \mathbf{A} is symmetric and positive definite ($\mathbf{A} > \mathbf{0}$) because it is the sum of a real, symmetric, positive semi-definite matrix, $\Phi[\mathbf{Y}_i(k), \mathbf{Y}_i(k)]$, and a real, symmetric, positive definite matrix, $\sigma_v^2 \mathbf{I}_{2k}$. Then, $\mathbf{A}^{-1} > \mathbf{0}$ and $\mathbf{B}^T \mathbf{A}^{-1} \mathbf{B} > \mathbf{0}_{2 \times 2}$, where \mathbf{B} is defined in (28), and $>$ denotes element-wise inequality. From (31), the following also holds,

$$\mathbf{Q} \leq \Psi[\mathbf{x}_j(k+1), \mathbf{x}_j(k+1)] + \sigma_v^2 \mathbf{I}_2 \quad (35)$$

and since the covariance matrix function is stationary, $\Psi[\mathbf{x}_j(k+1), \mathbf{x}_j(k+1)] = \Psi(\mathbf{0}, \mathbf{0}) = \Psi_0$ for any $\mathbf{x}_j \in \mathcal{W}$, and thus $\mathbf{Q} \leq \Psi_0 + \sigma_v^2 \mathbf{I}_2$.

Since $\Phi[\mathbf{Y}_i(k), \mathbf{Y}_i(k)]$ is real, symmetric, and positive semi-definite, there exists an eigenvalue decomposition, $\Phi[\mathbf{Y}_i(k), \mathbf{Y}_i(k)] = \mathbf{U} \boldsymbol{\Lambda} \mathbf{U}^{-1}$, with orthogonal eigenvectors, where $\boldsymbol{\Lambda}$ is a diagonal matrix obtained by placing the k eigenvalues of $\Phi[\cdot]$ on the diagonal, and \mathbf{U} is a $k \times k$ matrix whose columns are the eigenvectors of $\Phi[\cdot]$, i.e.,

$$\boldsymbol{\Lambda} \triangleq \text{diag}[\lambda_1 \cdots \lambda_k] \quad \text{and} \quad \mathbf{U} \triangleq [\mathbf{u}_1 \cdots \mathbf{u}_k]^T \quad (36)$$

and, thus, $\mathbf{U}^T \mathbf{U} = \mathbf{I}$. Since the k th column of $\mathbf{Y}_i(k)$ is equal to $\mathbf{x}_j(k)$, the matrix \mathbf{B} defined in (28) can be written as, $\mathbf{B} = \mathbf{U} \boldsymbol{\Lambda} \mathbf{u}_k$. By substituting \mathbf{B} into in (31), the matrix \mathbf{Q} can be written as,

$$\mathbf{Q} = \Psi_0 + \left[\sigma_v^2 - \sum_{\ell=1}^k \frac{\lambda_\ell}{\lambda_\ell + \sigma_v^2} \lambda_\ell (\mathbf{U}_{(\ell,k)})^2 \right] \mathbf{I}_2 \quad (37)$$

where λ_ℓ is the ℓ th eigenvalue of $\Phi[\mathbf{Y}_i(k), \mathbf{Y}_i(k)]$, and $\mathbf{U}_{(\ell,k)}$ denotes the element in the ℓ th row and k th column of \mathbf{U} . Because $\Phi[\cdot]$ is symmetric and positive semi-definite, $\lambda_\ell \geq 0$ for all ℓ , and, since $\mathbf{u}_k^T \boldsymbol{\Lambda} \mathbf{u}_k = 1$, it follows that $\sum_{\ell=1}^k \lambda_\ell (\mathbf{U}_{(\ell,k)})^2 = 1$, which can be substituted into (37) to show that, $\mathbf{Q} \geq \Psi_0 + \left(\sigma_v^2 - \frac{\max_{\ell} \{\lambda_\ell\}}{\max_{\ell} \{\lambda_\ell\} + \sigma_v^2} \right) \mathbf{I}_2$. Because the trace of a real, symmetric matrix equals the sum of its eigenvalues, and that the diagonal blocks of $\Phi[\mathbf{Y}_i(k), \mathbf{Y}_i(k)]$ are equal to $\Psi(\mathbf{y}_l, \mathbf{y}_l) = \Psi_0$, then $\sum_{\ell=1}^k \lambda_\ell = k \text{tr}(\Psi_0)$. Furthermore, $\max_{\ell} \{\lambda_\ell\} \leq k \text{tr}(\Psi_0)$, and thus,

$$\mathbf{Q} \geq \Psi_0 + \left[\frac{\sigma_v^2}{k \text{tr}(\Psi_0) + \sigma_v^2} \right] \mathbf{I}_2 \geq \Psi_0 > \mathbf{0} \quad (38)$$

which completes the proof. \square

The above result is used in the following lemma to establish a matrix inequality on consecutive covariance matrices $\Sigma_{i,k}$ and $\Sigma_{i,k+1}$.

Lemma 3. Under the assumptions in Proposition 1, two consecutive covariance matrices $\Sigma_{i,k}$ and $\Sigma_{i,k+1}$, defined according to (15), obey $\mathbf{0} < \Sigma_{i,k+1} \leq \Sigma_{i,k}$.

Proof. Under the assumptions in Proposition 1, $\Sigma_{i,k}, \Sigma_{i,k+1} > \mathbf{0}$ because they are Gaussian covariance matrices. From the block

matrix inversion of $\Sigma_{i,k+1}$ the difference between two consecutive covariance matrices can be written as

$$\Sigma_{i,k} - \Sigma_{i,k+1} = \mathbf{R}\mathbf{Q}^{-1}\mathbf{R}^T. \quad (39)$$

From Lemma 2, $\Psi_0 \leq \mathbf{Q}$ and, since \mathbf{Q} is a diagonal positive-definite matrix, it follows that \mathbf{Q}^{-1} is diagonal and positive definite. Then, there exists a diagonal positive-definite matrix, \mathbf{V} , such that $\mathbf{Q}^{-1} = \mathbf{V}\mathbf{V}^T$ and

$$\Sigma_{i,k} - \Sigma_{i,k+1} = \mathbf{R}\mathbf{V}\mathbf{V}^T\mathbf{R}^T = (\mathbf{R}\mathbf{V})(\mathbf{R}\mathbf{V})^T \succeq \mathbf{0}. \quad \square$$

The third and final lemma provides a bound on the trace of the quadratic form $\mathbf{R}^T\mathbf{R}$ that is later used to show that the Monte Carlo integration (33) is an unbiased estimator of the DPGP-EKLD in (24).

Lemma 4. *Under the assumptions in Proposition 1, the cross-covariance matrix \mathbf{R} , obtained from the target velocity and position estimates at the collocation points, and defined in (30), obeys the inequality $0 \leq \text{tr}(\mathbf{R}\mathbf{R}^T) \leq 4k[(\Psi_0) + 2\sigma_v^2]$, where $\Psi_0 = \Psi(\mathbf{0}, \mathbf{0})$ is a constant matrix obtained by evaluating the stationary covariance matrix function at zero, $\sigma_v \in \mathbb{R}$ is the standard deviation of the measurement noise, and $k \in \mathbb{Z}_+$ is the time index.*

Proof. Since $\mathbf{R}^T\mathbf{R}$ is a positive semi-definite matrix, it has positive or zero eigenvalues, and $\text{tr}(\mathbf{R}\mathbf{R}^T) \geq 0$. From GP regression, the joint probability distribution of a vector $[\mathbf{v}_i^T \ \mathbf{v}_j(k+1)^T \ \mathbf{f}_i(\mathbf{y}_1)^T \ \mathbf{f}_i(\mathbf{y}_2)^T \ \dots]^T$, comprised of target velocities at the collocation points (10) and at the measured target positions, is a multivariate Gaussian distribution with covariance matrix,

$$\begin{bmatrix} \Phi(\xi, \xi) + \sigma_v^2 \mathbf{I}_{2L} & \Phi[\xi, \mathbf{x}_j(k+1)] & \Phi[\xi, \mathbf{Y}_i(k)] \\ \Phi[\mathbf{x}_j(k+1), \xi] & \mathbf{D} & \mathbf{B}^T \\ \Phi[\mathbf{Y}_i(k), \xi] & \mathbf{B} & \mathbf{A} \end{bmatrix}$$

where $\mathbf{A}, \mathbf{B}, \mathbf{D}$ are defined in (27)–(29). Then, the conditional marginal distribution of the vector $[\mathbf{v}_i^T \ \mathbf{v}_j(k+1)^T]^T$, given the vector of target position measurements $\mathbf{Y}_i(k)$, is a multivariate Gaussian distribution with covariance matrix $\begin{bmatrix} \Sigma_{i,k} & \mathbf{R} \\ \mathbf{R}^T & \mathbf{Q} \end{bmatrix}$, where \mathbf{R} and \mathbf{Q} are defined in (30)–(31). Because the covariance matrix is symmetric positive definite, the off-diagonal elements of \mathbf{R} are smaller than the corresponding diagonal elements of \mathbf{Q} , or $\mathbf{R}_{(i,j)} < \mathbf{Q}_{(j,j)} < \text{tr}(\mathbf{Q})$, for all i, j . From Lemma 2, $\mathbf{Q} \leq \Phi_0 + \sigma_v^2 \mathbf{I}_2$, therefore it follows that $\mathbf{R}_{(i,j)} < \text{tr}(\Psi_0) + 2\sigma_v^2$ and, from the properties of quadratic forms, $\text{tr}(\mathbf{R}\mathbf{R}^T) = \sum_{ij} |\mathbf{R}_{(i,j)}|^2 < 4k[\text{tr}(\Psi_0) + 2\sigma_v^2]$, completing the proof. \square

We are now ready to prove the following theorem on the properties of the DPGP-EKLD approximation:

Theorem 5. *Under the assumptions in Proposition 1, the Monte Carlo (MC) integration (33) is an unbiased estimator of the DPGP-EKLD (24), and the variance of the difference between (24) and (33) decreases linearly with the number of samples, S , that are drawn independently and identically from the target state distribution (22).*

Proof. From the linearity of the expectation operation, the expected value of the DPGP-EKLD obtained via Monte Carlo integration in (33) is given by

$$\begin{aligned} \bar{D} &= \mathbb{E}_{\xi^{(l)}} \{ \hat{D}(\mathbf{v}; \mathbf{m}(k+1)) \} \\ &= \sum_{j=1}^N \sum_{i=1}^M \frac{w_{ij}}{S} \sum_{s=1}^S \mathbb{E}_{\xi^{(l)}} \{ h_i(\xi^{(l)}) \mathbf{1}_{\mathcal{S}(k+1)}(\xi^{(l)}) \}. \end{aligned} \quad (40)$$

Since MC samples are iid, the following holds,

$$\begin{aligned} \mathbb{E}_{\xi^{(l)}} \{ h_i(\xi^{(l)}) \mathbf{1}_{\mathcal{S}(k+1)}(\xi^{(l)}) \} \\ = \mathbb{E}_{\xi^{(l)}} \{ h_i[\mathbf{x}_j(k+1)] \mathbf{1}_{\mathcal{S}(k+1)}[\mathbf{x}_j(k+1)] \} \end{aligned} \quad (41)$$

for all $l = 1, \dots, S$, where $h_i[\mathbf{x}_j(k+1)]$ is defined in (26), and thus the estimator (33) is unbiased. Also, the variance of the DPGP-EKLD in (40) can be written as,

$$\begin{aligned} \text{var}(\hat{D}) &= \mathbb{E}\{(\hat{D} - \bar{D})^2\} \\ &= \sum_{j=1}^N \sum_{i=1}^M w_{ij} \frac{1}{S^2} \sum_{l=1}^S \mathbb{E} \left\{ \left[h_i(\xi^{(l)}) \mathbf{1}_{\mathcal{S}(k+1)}(\xi^{(l)}) \right. \right. \\ &\quad \left. \left. - \mathbb{E} \left\{ h_i[\mathbf{x}_j(k+1)] \mathbf{1}_{\mathcal{S}(k+1)}[\mathbf{x}_j(k+1)] \right\} \right]^2 \right\} \\ &= \frac{1}{S} \text{var} \left\{ h_i[\mathbf{x}_j(k+1)] \mathbf{1}_{\mathcal{S}(k+1)}[\mathbf{x}_j(k+1)] \right\} \end{aligned} \quad (42)$$

where it can be proven that $\text{var} \{ h_i[\mathbf{x}_j(k+1)] \mathbf{1}_{\mathcal{S}(k+1)}[\mathbf{x}_j(k+1)] \}$ is a finite constant because $h_i[\cdot]$ is finite-valued as follows.

As a first step, we prove that $h_i[\cdot]$ is greater than or equal to zero. From (25), $h_i[\cdot]$ is the integral of the KL divergence weighted by a Gaussian distribution. Since the KL divergence is always greater than or equal to zero, and $p(\mathbf{z}_j(k+1) | \mathbf{x}_j(k+1)) \geq 0$, it also follows that $h_i[\cdot] \geq 0$. As a second step, we prove that $h_i[\cdot]$ is finite by showing that every term of $h_i[\cdot]$ is finite. The first term in (26) is finite-valued because from Lemma 3,

$$\text{tr}(\Sigma_{i,k}^{-1} \Sigma_{i,k+1}) = \text{tr}(\mathbf{I}_{2L} - \Sigma_{i,k} \mathbf{R} \mathbf{Q}^{-1} \mathbf{R}^T) \leq 2L \quad (43)$$

where L is the number of collocation points in \mathcal{W} . The second term in (26) is also finite because from Lemma 3 $|\Sigma_{i,k}| > 0$ and $|\Sigma_{i,k+1}| > 0$, since the two covariance matrices are positive definite. Since $\Sigma_{i,k}$ and $\Sigma_{i,k+1}$ are Hermitian and $\Sigma_{i,k} \succeq \Sigma_{i,k+1}$, then $|\Sigma_{i,k}| > |\Sigma_{i,k+1}|$ and $0 < |\Sigma_{i,k+1}|/|\Sigma_{i,k}| < 1$, such that $0 < -\ln(|\Sigma_{i,k+1}|/|\Sigma_{i,k}|) < \infty$. Since the trace of a product of matrices is invariant under cyclic permutations, and the trace of the product of positive semi-definite matrices is less than or equal to the product of the individual traces (Bernstein, 2005), it follows that,

$$\begin{aligned} \text{tr}(\mathbf{Q}^{-1} \mathbf{R}^T \Sigma_{i,k}^{-1} \mathbf{R} \mathbf{Q}^{-1}) \sigma_v^2 &= \text{tr}(\Sigma_{i,k}^{-1}) [\text{tr}(\mathbf{Q}^{-1})]^2 \\ &\times \text{tr}(\mathbf{R}\mathbf{R}^T) \sigma_v^2 \leq 4\text{tr}(\Sigma_{i,k}^{-1}) [4k\text{tr}(\Phi_0) + 2\sigma_v^2] \sigma_v^2 \end{aligned} \quad (44)$$

where the last inequality is the result of Lemmas 2–4. Now that every term in (26) has been proven finite, $h_i[\cdot]$ is finite-valued for all $\mathbf{x}_j(k+1) \in \mathbb{R}^2$. Also from (34), the variance in (42) is a definite integral of $h_i[\cdot]$ over $\mathcal{S}(k+1)$, and thus a finite constant. Then, the error variance of the MC estimator in (33) is proportional to $1/S$ (Caflich, 1998). \square

5. DPGP sensor planning

The DPGP-EKLD approximation in (33) is used to plan the sensor measurements in the target modeling problem formulated in Section 2. The performance obtained by planning the motion of the sensor FoV, $\mathcal{S}(k+1)$, using a DPGP-EKLD maximizing algorithm is compared to an algorithm based on DPGP mutual information (MI), derived by the approach in Section 4, a particle filter (PF) algorithm (Wang et al., 2008), and a randomized method (Fulgenci et al., 2008). The workspace is $\mathcal{W} = [0, 10] \times [0, 10]$ m², and the camera FoV is a square of side $s = 0.5$ m² and variable position in \mathcal{W} . The sensor measurements are characterized by $\sigma_x = \sigma_v = 0.1$ (m/s) and $\Delta t = 0.3$ s. The unknown target ODEs are,

$$\begin{aligned} \mathbf{f}_1[\cdot] &= [-\sin([0.1 \ \pi/20]\mathbf{x}_j) \ \sin([\pi/12 \ 1]\mathbf{x}_j + \pi/4)]^T \\ \mathbf{f}_2[\cdot] &= [-\cos([0 \ \pi/8]\mathbf{x}_j) \ \sin([\pi/4 \ 0]\mathbf{x}_j)]^T \\ \mathbf{f}_3[\cdot] &= [-0.5 \ \sin([\pi/4 \ 0.3]\mathbf{x}_j)]^T \\ \mathbf{f}_4[\cdot] &= [-\cos([\pi/8 \ -0.5]\mathbf{x}_j) \ 1]^T \end{aligned} \quad (45)$$

such that $\mathcal{F} = \{\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3, \mathbf{f}_4\}$, and $\boldsymbol{\pi} = [0.25 \ 0.25 \ 0.25 \ 0.25]^T$.

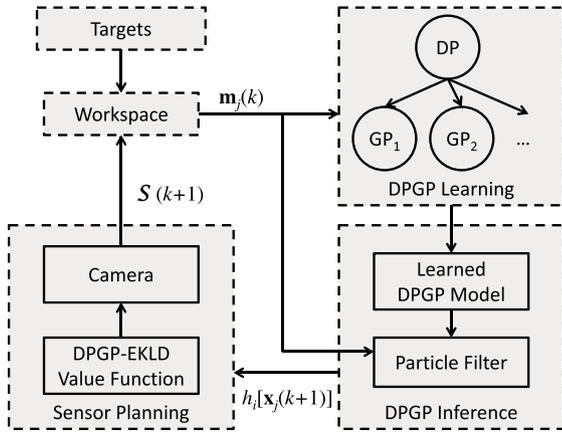


Fig. 1. Diagram of DPGP-EKLD sensor planning algorithms.

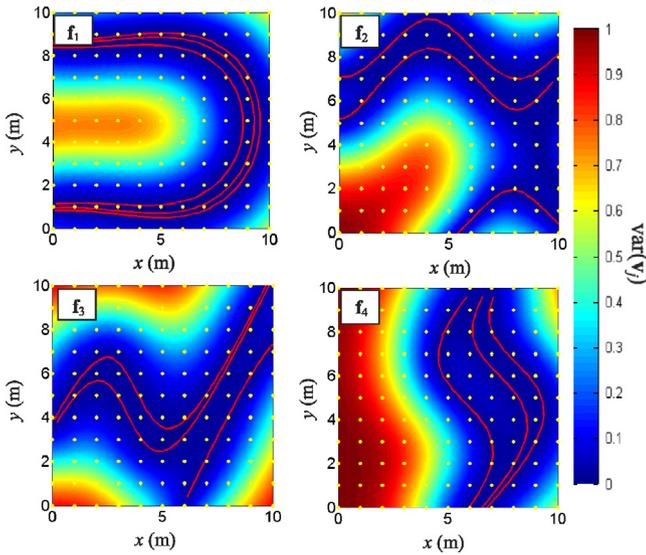


Fig. 2. Example target trajectories (red), collocation points (yellow) and velocity variance (contour plot) for the four velocity fields in (45). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The performance of the DPGP-EKLD algorithm is evaluated for $t_0 = 0$, $t_f = 300$ s. During this time, $N = 500$ targets enter and exit \mathcal{W} following a VF in \mathcal{F} selected randomly from $\text{Cat}(\pi)$. The target VFs and statistics $\{\mathcal{F}, \pi\}$ are learned from data, using $\Psi(\mathbf{x}_j, \mathbf{x}'_j) = \exp(-\|\mathbf{x}_j - \mathbf{x}'_j\|_2^2 / 2\ell^2) \mathbf{I}_2$ and $\ell = \sqrt{10}$ based on the size of \mathcal{W} . An example of target trajectories used to learn the DPGP model is shown in Fig. 2, superimposed on a contour plot of $\text{var}(\mathbf{v}_j)$. At the onset of the simulation, the DP concentration parameter is chosen as $\alpha = 1$. Subsequently, the DPGP model is updated at every time t_k , based on the complete database $Q(k)$. As schematized in Fig. 1, the target-VF associations $\mathcal{g}(k)$ are obtained from the DPGP using a particle filter (Wei et al., 2014). Based on the latest DPGP, the position of the camera FoV in \mathcal{W} is computed so as to maximize EKLD and obtain the target measurements \mathbf{m}_j with the highest information value.

From (3), the DPGP target model can be evaluated by comparing its estimate of target velocity, $\hat{\mathbf{v}}_j(k) = \theta_i(\mathbf{x}_j) = \mathbb{E}[\mathbf{f}_i(\mathbf{x}_j)]$, to the actual velocity $\mathbf{v}_j(k)$ such that, at t_k , the average root mean squared error (RMSE) is,

$$\varepsilon_k = \frac{1}{N} \sum_{j=1}^N \sum_{i=1}^M w_{ij} \sqrt{\frac{1}{K_j} \sum_{k=1}^{K_j} \|\mathbf{v}_j(k) - \hat{\mathbf{v}}_j(k)\|_2^2} \quad (46)$$

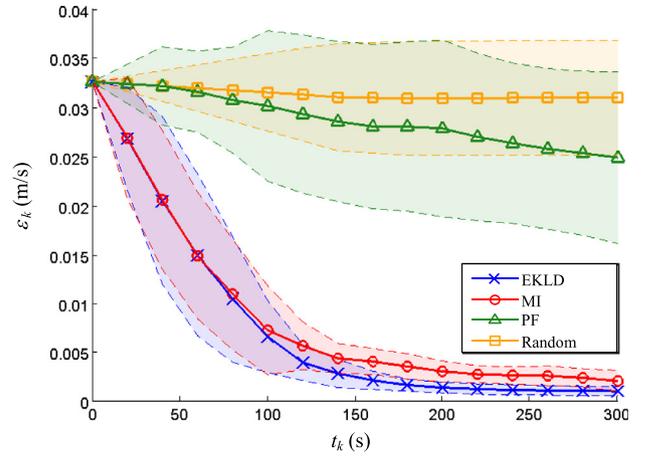


Fig. 3. Mean and (\pm) standard deviation of the VF RMSE in (46) resulting from EKLD, MI, PF, and random camera planning, for the case of medium prior information.

Table 1
Average RMSE of DPGP model.

Algorithms:	EKLD	MI	PF	Random
Low prior information	0.35%	0.79%	3.31%	5.16%
Medium prior information	0.11%	0.25%	2.46%	3.32%
High prior information	0.09%	0.21%	0.56%	1.34%

where K_j is the number of time intervals spent in \mathcal{W} by target j , and w_{ij} is the target-VF association probability in (19). The EKLD algorithm is compared to an MI algorithm that maximizes DPGP mutual information, a PF algorithm that places $\mathcal{S}(k+1)$ to track the nearest target, and a random algorithm that places $\mathcal{S}(k+1)$ by sampling a uniform multivariate distribution with support \mathcal{W} . Three cases with low, medium, and high prior information are considered, obtaining the results in Table 1. It can be seen that the EKLD information function is the most effective at improving RMSE, or reducing model error, over time.

The time history of RMSE and its variance for all four planning algorithms is plotted in Fig. 3, for the case of medium prior information. These results, obtained for a sample of fifty simulations performed under the same DPGP statistics, illustrate the performance robustness of the EKLD algorithm. The MI algorithm also displays good performance at the onset of learning, because the medium prior is approximately equally informative about all four VFs. However, as more target measurements of targets are obtained over time, the EKLD function is more effective because it prioritizes targets characterized by VFs with higher uncertainty. Finally, the results in Fig. 4 show that the posterior probabilities of target-VF association defined in (19), and obtained from the updated DPGP via PF, converge to their true values (dash-dotted lines) over time. The snapshots in Fig. 4(a)–(c) show that, by planning the camera movements via EKLD, the FoV is able to intersect PF clusters with high information value and, thus, rapidly improve their posterior probability distributions.

6. Conclusion

This paper presents an approach for deriving tractable information value functions for nonparametric DPGP mixture models. Although information theoretic functions are a common approach for representing information value in sensing and control problems, they are not directly applicable to DPGP models of stochastic processes because they are typically defined in terms of finite-dimensional probability distributions. The proposed approach is demonstrated by deriving efficient representations of KL

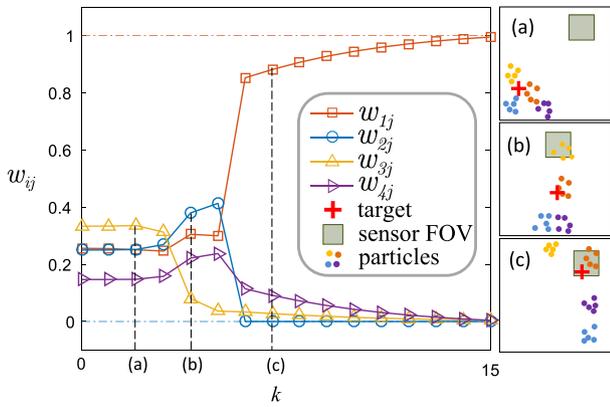


Fig. 4. Time history of target-VF posterior probabilities updated over time by the DPGP and PF algorithms in Fig. 1.

divergence in DPGP models via collocation methods and MC integration. In particular, the EKLD approximation, useful for planning and control of mobile sensors, is proven to be an unbiased estimator of the original information function, with an error variance that is guaranteed to decrease linearly with the number of MC samples. The effectiveness of the approach is demonstrated through a sensor planning problem in which the motion of the camera FoV is planned by optimizing the DPGP-EKLD value function over time. The results show that the EKLD sensor planning algorithm outperforms algorithms based on DPGP mutual information, particle filter, and randomized methods.

Appendix. Proof of Eq. (16)

Proof. When $G_j(k + 1) = i$,

$$p(v_i|Q(k + 1)) = p(v_i|Q(k + 1)) \prod_{1 \leq l \leq M, l \neq i} p(v_l|Q(k)) \tag{A.1}$$

since $v_i = v_i(\xi)$ and $v_l = v_l(\xi)$ are conditional independent given $Q(k + 1)$. Substituting (A.1) into (12), the conditional KL divergence can be written as

$$\begin{aligned} D(v; \mathbf{m}_j(k + 1)|Q(k)) &= \int_{\mathbb{R}^{2M}} \ln \left\{ \frac{p(v_i|Q(k + 1)) \prod_{1 \leq l \leq M, l \neq i} p(v_l(\xi)|Q(k))}{\prod_{\ell=1}^{2M} p(v_\ell(\xi)|Q(k))} \right\} \\ &\quad \times p(v|Q(k + 1)) dv \\ &= \int_{\mathbb{R}^{2M}} \ln \left\{ \frac{p(v_i|Q(k + 1))}{p(v_i|Q(k))} \right\} p(v_i|Q(k + 1)) dv_i \\ &= D(v_i; \mathbf{m}_j(k + 1)|Q(k)). \quad \square \end{aligned} \tag{A.2}$$

References

Bar-Shalom, Y., & Blair, W. D. (2000). *Multitarget-multisensor tracking: Applications and advances, vol. III*. Artech House.
 Bar-Shalom, Y., & Tse, E. (1975). Tracking in a cluttered environment with probabilistic data association. *Automatica*, 11(5), 451–460.
 Bernstein, D. S. (2005). *Matrix mathematics*. Princeton, NJ: Princeton University Press.
 Bogachev, V. I., & Ruas, M. A. S. (2007). *Measure theory, vol. 1*. Springer.
 Caffisch, R. E. (1998). Monte Carlo and Quasi-Monte Carlo methods. *Acta Numerica*, 7, 1–49.
 Cover, T., & Thomas, J. (1991). *Elements of information theory*. New York, NY: Wiley-Interscience.

Fulgenzi, C., Tay, C., Spalanzani, A., & Laugier, C. (2008). Probabilistic navigation in dynamic environment using rapidly-exploring random trees and Gaussian processes. In *Intelligent robots and systems, 2008. IROS 2008. IEEE/RSJ international conference on* (pp. 1056–1062). IEEE.
 Hensman, J., Rattray, M., & Lawrence, N. D. (2015). Fast nonparametric clustering of structured time-series. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2), 383–393.
 Joseph, J., Doshi-Velez, F., Huang, A. S., & Roy, N. (2011). A Bayesian nonparametric approach to modeling motion patterns. *Autonomous Robots*, 31(4), 383–400.
 Kastella, K. (1997). Discrimination gain to optimize detection and classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 27(1), 112–116.
 Krause, A., Singh, A., & Guestrin, C. (2008). Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research (JMLR)*, 9, 235–284.
 Lu, W., Zhang, G., & Ferrari, S. (2014). An information potential approach to integrated sensor path planning and control. *IEEE Transactions on Robotics*, 30(4), 919–934.
 Nobile, F., Tempone, R., & Webster, C. G. (2008). A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 46(5), 2309–2345.
 Ny, J. L., & Pappas, G. J. (2009). On trajectory optimization for active sensing in Gaussian process models. In *Proc. of the IEEE conference on decision and control* (pp. 6286–6292). Shanghai, China, Dec.
 Pattipati, K. R., Deb, S., Bar-Shalom, Y., & Washburn, R., Jr. (1992). A new relaxation algorithm and passive sensor data association. *IEEE Transactions on Automatic Control*, 37(2), 198–213.
 Press, W. H. (2007). *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press.
 Rasmussen, C. E. (1999). The infinite Gaussian mixture model. In *Proc. of advances in neural information processing systems, NIPS: Vol. 12* (pp. 554–560). USA: Denver, CO, Dec.
 Rasmussen, C. E., & Williams, C. (2006). *Gaussian processes for machine learning*. MIT Press.
 Ross, J., & Dy, J. (2013). Nonparametric mixture of Gaussian processes with constraints. In *proceedings of the 30th international conference on machine learning, ICML-13* (pp. 1346–1354).
 Teh, Y. W., Jordan, M. I., Beal, M. J., & Blei, D. M. (2006). Hierarchical Dirichlet processes. *Journal of the American Statistical Association*, 101(476).
 Wang, J., Yin, Y., & Hong, M. (2008). Multiple human tracking using particle filter with Gaussian process dynamical model. *EURASIP Journal on Image and Video Processing*.
 Wei, H., Lu, W., & Ferrari, S. (2012). An information value function for nonparametric Gaussian processes. In *proc. neural information processing systems conference, Lake Tahoe, NV*.
 Wei, H., Lu, W., Zhu, P., Ferrari, S., Klein, R. H., Omidshafiei, S., & How, J. P. (2014). Camera control for learning nonlinear target dynamics via Bayesian nonparametric Dirichlet-Process Gaussian-Process (DP-GP) models. In *Intelligent robots and systems, IEEE/RSJ international conference on* (pp. 95–102). IEEE.
 Wei, H., Lu, W., Zhu, P., Huang, G., Leonard, J., & Ferrari, S. (2014). Optimized visibility motion planning for target tracking and localization. In *Intelligent robots and systems, 2014 IEEE/RSJ international conference on* (pp. 76–82). IEEE.
 Wei, H., Ross, W., Varisco, S., Krief, P., & Ferrari, S. (2013). Modeling of human driver behavior via receding horizon and artificial neural network controllers. In *decision and control, IEEE annual conference on* (pp. 6778–6785). Florence, Italy, Dec.
 Zhang, G., Ferrari, S., & Cai, C. (2012). A comparison of information functions and search strategies for sensor planning in target classification. *IEEE Transactions on Systems, Man, and Cybernetics, - Part B*, 42(1), 2–16.



Hongchuan Wei earned his B.S. degree in Vehicle Engineering from Tsinghua University, China in 2010 and is pursuing a Ph.D. degree at Duke University. His research interest includes sensor network deployment, which consists of optimally placing a set of sensors in a region of interest (ROI) such that the probability of detection of targets traveling through the ROI is maximized. The probability of detection is a function of the target movement and sensors field of view. He is also interested in utilizing information theoretic functions to acquire optimal sensor controls in order to learn the behaviors of targets described by non-parametric Bayesian models.



Wenjie Lu received the Ph.D. (2014) and M.S. (2011) degree in mechanical engineering from Duke University, Durham, NC, USA, and the B.S. (2009) degree in mechatronic engineering from Zhejiang University, Hangzhou, Zhejiang, China. He is currently a Postdoctoral Fellow with CAS at University of Technology Sydney, Australia. His research interests include intelligent underwater vehicle manipulator systems, heterogeneous mobile sensors, approximate dynamic programming, hybrid systems, information theory, and nonparametric Bayesian models.



Pingping Zhu received the B.S. degree in electronics and information engineering and the M.S. degree in the institute for pattern recognition and artificial intelligence from Huazhong University of Science and Technology, and M.S. and Ph.D. degrees in electrical and computer engineering from University of Florida. He is currently a research associate at the department of mechanical and aerospace engineering in Cornell University. Prior to that, he was a postdoctoral associate at the department of mechanical engineering and material science in Duke university. His research interests include approx-

imate dynamic programming, optimal control of mobile sensor networks, signal processing, machine learning, and neural networks. He is a member of the IEEE.



Silvia Ferrari is a Professor of MAE at Cornell University. Prior to that, she was Professor of Engineering and Computer Science, and Founder and Director of the NSF Integrative Graduate Education and Research Traineeship (IGERT) and Fellowship program on Wireless Intelligent Sensor Networks (WiSeNet) at Duke University. She is the Director of the Laboratory for Intelligent Systems and Controls (LISC), and her principal research interests include robust adaptive control of aircraft, learning and approximate dynamic programming, and optimal control of mobile sensor networks. She received the B.S. degree

from Embry–Riddle Aeronautical University and the M.A. and Ph.D. degrees from Princeton University. She is a senior member of the IEEE, and a member of ASME, SPIE, and AIAA. She is the recipient of the ONR young investigator award (2004), the NSF CAREER award (2005), and the Presidential Early Career Award for Scientists and Engineers (PECASE) award (2006).



Miao Liu is a postdoctoral associate in the Laboratory of Information and Decision System (LIDS) at MIT, where he works on scalable Bayesian nonparametric methods for solving decentralized multiagent learning and planning problems. He received his Ph.D. in Electrical and Computer Engineering from Duke University in 2014. He received his B.S. and M.S. in Electronics and Information Engineering from Huazhong University of Science and Technology, in Wuhan, China in 2005 and 2007, respectively. His research interests include statistical signal and image processing, machine learning and artificial intelligence.



Robert H. Klein is currently a Security Researcher at Microsoft. His primary role is to leverage high volume telemetry feeds to automate discovery of targeted computer network attacks. Previously he was a data scientist and software engineer working on the BluVector security appliance (2014–2016), and a Systems Engineer in Northrop Grumman's Future Technical Leaders Program (2014). He received a Bachelor's Degree in Mechanical and Aerospace Engineering from Princeton University in 2012, and his Master's Degree in Aeronautics and Astronautics from MIT in 2014, where he studied Machine Learning in the Aerospace Controls Laboratory. Bob is a National Science Foundation Graduate Research Fellow, and a recipient of the Jack and Vickie Kerrebrock Fellowship in Aerospace Engineering.



Shayegan Omidshafiei is a Ph.D. student at the Laboratory for Information and Decision Systems (LIDS) and Aerospace Controls Laboratory (ACL) at MIT, Cambridge, MA. He received the B.A.Sc. degree from the University of Toronto in 2012 and the S.M. degree from MIT in Aeronautics and Astronautics under supervision of Professor Jonathan How. His research interests include multi-agent planning under uncertainty, autonomous vehicles, reinforcement learning, and data visualization.



Jonathan P. How is the Richard C. Maclaurin Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology. He received a B.A.Sc. from the University of Toronto in 1987 and his S.M. and Ph.D. in Aeronautics and Astronautics from MIT in 1990 and 1993, respectively. He then studied for two years at MIT as a postdoctoral associate for the Middeck Active Control Experiment (MACE) that flew onboard the Space Shuttle Endeavour in March 1995. Prior to joining MIT in 2000, he was an Assistant Professor in the Department of Aeronautics and Astronautics at Stanford University.

He is the Editor-in-chief of the IEEE Control Systems Magazine and an Associate Editor for the AIAA Journal of Aerospace Information Systems. Professor How was the recipient of the 2002 Institute of Navigation Burka Award, a Boeing Special Invention award in 2008, the IFAC Automatica award for best applications paper in 2011, the AeroLion Technologies Outstanding Paper Award for the Journal Unmanned Systems in 2015, won the IEEE Control Systems Society Video Clip Contest in 2015, and received the AIAA Best Paper in Conference Awards in 2011, 2012, and 2013. He is a Fellow of AIAA and a senior member of IEEE.