2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

# Categorical Soft Data Fusion via Variational Bayesian Importance Sampling, with Applications to Cooperative Search

Nisar Ahmed*, Eric Sample*, Ken Ho+, Tauhira Hoossainy* and Mark Campbell*

*Abstract*— **This paper considers Bayesian data fusion with categorical 'soft sensor' information, such as human input in cooperative multi-agent search applications. Previous work studied variational Bayesian (VB) hybrid data fusion, which produces optimistic posterior covariance estimates and requires simple Gaussian priors with softmax likelihoods. Here, a new hybrid fusion procedure, known as variational Bayesian importance sampling (VBIS), is introduced to combine the strengths of VB and fast Monte Carlo methods to produce more reliable Gaussian posterior approximations for Gaussian priors and softmax likelihoods. VBIS is then generalized to problems involving complex Gaussian mixture priors and multimodal softmax observation models to obtain reliable Gaussian mixture posterior approximations. The utility and accuracy of the VBIS fusion method is demonstrated on a multitarget search problem for a real cooperative human-automaton team.**

## I. INTRODUCTION

Control and estimation problems involving cooperative teams of multiple heterogeneous agents have received much attention in the past decade, due to their wide applicability to areas such as defense [1] and search and rescue [2]. While control and sensing capabilities for autonomous unmanned vehicles are always improving for such applications, human agents still serve key roles both as supervisory operators and as information-sharing agents [3]. Although human information provision is often geared towards classifying objects and abstract behaviors, less appreciated is the fact that humans can provide 'soft' information about continuous quantities of interest. Though not as precise as conventional 'hard' sensor information, soft information has been shown to be useful for real-world data fusion problems such as target tracking [4] and robot localization [5].

Refs. [3], [6], [7] explored incorporation of human information sources in the context of surveillance and search-and-rescue applications. However, few formal methods have been proposed for 'low-level' human information fusion in the Bayesian framework, which has become increasingly popular for such applications. Ref. [3] developed a Bayesian method for fusing a human's continuous range/bearing estimates to tracked objects by modeling the 'human sensor' via linear regression models. Ref. [7] considered grid-based Bayesian fusion of human detection likelihoods for a distributed 2D search problem, where human field of view is modeled as a soft binary sensor likelihood. While humans tend to categorize continuous data as a convenient means for sharing and processing complex information, efficient Bayesian soft information fusion remains challenging, given the potentially complex dependencies between continuous states and non-binary categorical observations. Some approximations for hybrid fusion problems are considered in [8], although these have some drawbacks that limit general applicability.

This paper describes how to efficiently fuse contextual multi-categorical information about continuous states in a formal Bayesian framework that is applicable to general fusion scenarios. There are three contributions. Firstly, a new variational Bayesian importance sampling (VBIS) procedure is presented for obtaining reliable hybrid data fusion posteriors in the baseline case of Gaussian priors with softmax observation likelihoods, which was studied in [8]. Secondly, VBIS is extended to the more general case of Gaussian mixture priors and multimodal softmax observation models. This leads to a Gaussian mixture respresentation for fusion updates that is ideal for complex recursive Bayesian estimation scenarios. Finally, the utility and accuracy of the VBIS fusion method is demonstrated through a multi-target search experiment with a real cooperative human-robot team. Comparisons to other methods for categorical human information fusion show that VBIS provides a more reliable representation of complex hybrid fusion posteriors.

## II. HYBRID BAYESIAN FUSION FOR CONTINUOUS STATES AND CATEGORICAL DATA

### A. General hybrid Bayesian fusion formulation

Suppose that $X \in \mathbb{R}^n$ is a random state variable with prior pdf $p(X)$ and that $D$ is an $m$-valued discrete random variable representing a categorical observation of $X$ with likelihood function $P(D = j|X)$, where $j \in \{1, ..., m\}$ and $\sum_{j=1}^{m} P(D = j|X = x) = 1$ for any realization $X = x$. To perform hybrid data fusion and update the uncertainty in $X$ given a new observation $D = j$, the posterior pdf $p(X|D = j)$ can be found via Bayes' rule,

$$p(X|D=j) = \frac{p(X, D=j)}{P(D=j)} = \frac{p(X)P(D=j|X)}{\int p(X)P(D=j|X)dX}. \quad (1)$$

In the context of heterogeneous multisensor fusion with human agents, $X$ can be any set of continuous variables of variable of interest (e.g. the 3D location of an object). $D$ models a human agent's soft categorical observation of $X$ such as 'The object is nearby Building A', where 'nearby' can be modeled as an element of some mutually exclusive set of contextual soft distance measures to which $X$ maps (e.g. $D \in \{$'Next To', 'Nearby', 'Far From', 'Very Far From'$\}$ as a function of the object's distance to Building A).

* Autonomous Systems Laboratory, Cornell University, Ithaca NY, 14853, USA, email: (nra6,ems42,ths432,mc288)@cornell.edu
+ Australian Center for Field Robotics, University of Sydney, Sydney, NSW 2006 Australia, email: k.ho@acfr.usyd.edu.au

### B. VB Fusion with Gaussian priors with softmax likelihoods

Assume that $X$ has a Gaussian prior $p(X) = \mathcal{N}(\mu, \Sigma)$ with mean $\mu \in \mathbb{R}^n$ and covariance $\Sigma \in \mathbb{R}^{n \times n}$. Furthermore, assume the likelihood function $P(D = j|X = x)$ is described by the softmax function,

$$P(D = j|X = x) = \frac{e^{w_j^T x + b_j}}{\sum_{k=1}^{m} e^{w_k^T x + b_k}}, \tag{2}$$

where $w_k \in \mathbb{R}^n$ is a vector weight and $b_k$ is a scalar bias for discrete class labels $j, k \in \{1, ..., m\}$. $P(D|X)$ can be identified from training data via standard learning algorithms [9], while $p(X)$ can be treated either as an initial belief or as the output of another Bayesian estimator (e.g. an EKF for continuous sensor data fusion [10]). Substituting $p(X)$ and (2) into (1) yields the exact posterior

$$p(X|D = j) = \frac{1}{C} \left|2\pi\Sigma^{-1}\right| e^{-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)} \frac{e^{w_j^T x + b_j}}{\sum_{c=1}^{m} e^{w_c^T x + b_c}}. \tag{3}$$

This posterior is not closed-form, since $C = P(D = j) = \int p(X)P(D = j|X)dX$ is analytically intractable. To cope with this problem, [8] exploits the fact that $p(X|D = j)$ is unimodal and can be well-approximated by a Gaussian,

$$p(X|D = j) \approx \hat{p}(X|D = j) = \frac{\hat{p}(X, D = j)}{\hat{C}} = \mathcal{N}(\hat{\mu}, \hat{\Sigma}), \tag{4}$$

where, from (1),

$$p(X, D = j) \approx \ \hat{p}(X, D = j) = p(X)\ f(D = j, X, \alpha, \xi) \tag{5}$$

$$C \approx \ \hat{C} = \int \hat{p}(X, D = j)\ dX = \hat{P}(D = j). \tag{6}$$

Here, the softmax function (2) is replaced by an unnormalized Gaussian lower bound $f(D = j, X, \alpha, \xi)$ derived from [11], so that the joint pdf is approximated by an unnormalized (but analytically integrable) Gaussian in (5). This is known as a *variational Bayes* (VB) approximation, since the size/shape variables $\alpha$ and $\xi$ are selected to maximize $\hat{C}$ in (6) as a lower bound to $C$. The unknown coupled parameters $\xi$, $\alpha$, $\hat{\mu}$ and $\hat{\Sigma}$ are estimated through a closed-form expectation maximization (EM) algorithm that converges monotonically to a unique solution. See [8] for complete details on the VB algorithm.

As shown in [8] and in Section III-A here, VB fusion provides good (though slightly biased) $\hat{\mu}$ estimates and optimistic $\hat{\Sigma}$ estimates, as is typical of VB approximations [9]. While reliability in $\hat{\mu}$ provides robustness to 'surprising observations' or incorrect prior information, optimism in $\hat{\Sigma}$ can lead to severe estimation inconsistencies due to overconfidence [10]. Furthermore, the assumption of Gaussian priors with softmax likelihoods restricts VB fusion to cases where the posterior is expected to be unimodal and well-approximated by a Gaussian. However, in practical hybrid fusion applications, multimodal posteriors can arise from non-Gaussian priors or complex categorical observation likelihoods for which (2) is inadequate. These issues are addressed in the next section.

## III. VARIATIONAL BAYESIAN IMPORTANCE SAMPLING FOR HYBRID DATA FUSION

This section presents the new variational Bayesian importance sampling (VBIS) algorithm, which produces more reliable estimates for the hybrid fusion approximation (4) for Gaussian state priors and softmax likelihoods. VBIS is then generalized to other hybrid fusion problems involving more complex Gaussian mixture (GM) state priors and multimodal softmax (MMS) observation likelihoods. This leads to approximate hybrid fusion posteriors in the form of GM models, which are well-suited to applications with nonlinear/non-Gaussian process and observation models.

### A. VBIS for Gaussian priors with softmax likelihoods

The VBIS approximation strategically combines VB fusion's reliability in $\hat{\mu}$ with the flexibility of Monte Carlo importance sampling (IS) [12]. This leads to fast and reliable estimates in both $(\hat{\mu}, \hat{\Sigma})$ for (4), since IS can estimate any statistical moment of (3) if enough samples are drawn from an 'importance distribution' $q(X, D = j)$ that is 'close' to (3). As the example below shows, IS is unreliable if $q$ is chosen poorly, e.g. if $q = p(X)$ is always used when the prior information is inconsistent with observations. Hence, by first using $\hat{\mu}$ from VB to define a suitable $q$ that is close to (3), IS can then be used to accurately re-estimate $\hat{\mu}$ and estimate $\hat{\Sigma}$. The detailed procedure for VBIS is as follows:

1) given $p(X)$ and $P(D = j|X)$, estimate $\hat{\mu}$ using VB (Sec. II-B and ref. [8])
2) set the importance sampling distribution to be
$$q(X, D = j) = \mathcal{N}(\hat{\mu}, \Sigma), \tag{7}$$
where $\Sigma$ is the covariance matrix of $p(X)$
3) draw $N_s$ samples $\{x_i\}_{i=1}^{N_s}$ from $q(X, D = j)$
4) compute importance weights $W_i$ for each sample $x_i$,
$$W_i \propto \frac{p(X = x_i)P(D = j|X = x_i)}{q(X = x_i, D = j)}, \tag{8}$$
and normalize so that $\sum_{i=1}^{N_s} W_i = 1$
5) re-estimate $\hat{\mu}$ and the posterior covariance $\hat{\Sigma}$,
$$\hat{\mu} = \sum_{i=1}^{N_s} W_i x_i, \quad \hat{\Sigma} = \sum_{i=1}^{N_s} W_i (x_i x_i^T) - \hat{\mu}\hat{\mu}^T \tag{9}$$

In theory, $q$ can be any pdf which is easy to sample and ensures proper coverage of (3) in high density areas. Eq. (7) is thus justified, since: (i) $\hat{\mu}$ ensures that $q$ is near a region of high probability mass, and (ii) it is expected that $(\Sigma - \hat{\Sigma})$ is positive definite, as conditioning the joint pdf (which is unimodal in $X$) on $D = j$ reduces the uncertainty in the (unimodal) posterior pdf [10]. Figure 1 demonstrates VBIS on two example 1D hybrid fusion scenarios, where $p(X)$ and $D$ agree well in (a) but disagree strongly in (b). Fusion results for VBIS (green, $N_s = 500$), stand-alone VB (red) and likelihood weighted IS (LWIS) (blue, $N_s = 500$) are shown for comparison to the true posterior (pink, numerical integration). LWIS (used in 'boostrap' particle filters [12]) uses $q = p(X)$ in (7), so (8) becomes $W_i \propto$

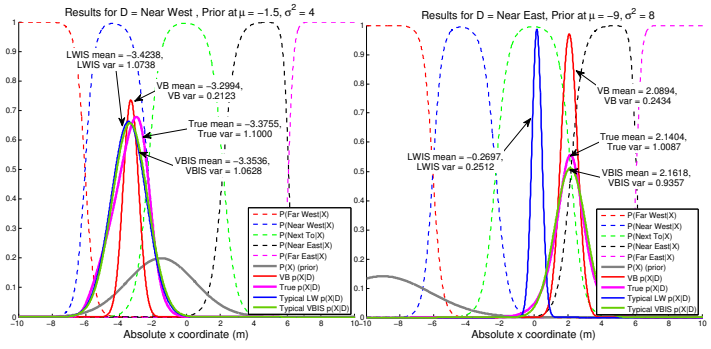


(a) (b)

Fig. 1. 1D fusion example: the position $X$ of an object has prior $p(X) = \mathcal{N}(\mu, \sigma)$ (gray) with $(\mu, \sigma^2)$ shown at the top. The softmax likelihood $P(D|X)$ has 5 discrete class labels describing the range/bearing of the object relative to $x = 0$; the likelihood contours for each class are shown as dashed lines and the observation in each case is shown at the top. Note that pdf plots are scaled for clarity.

$P(D = j|X)$. In each case, $\hat{\mu}_{VB}$ has a small bias and $\hat{\sigma}^2_{VB}$ is clearly optimistic, while VBIS produces a better Gaussian approximation to the true posterior. LWIS peforms erratically, as it is sensitive to the agreement between $p(X)$ and $D$. The VBIS computation time (10 ms in Matlab) is close to stand-alone VB, while LWIS is faster (1 ms). Repeating LWIS in case (b) with $N_s = 14,000$ samples gives 10 ms computation time and still less accurate results than VBIS, with $(\hat{\mu}_{LWIS}, \hat{\sigma}^2_{LWIS})$=(1.2895, 0.6390).

### B. VBIS for GM prior and MMS likelihood

*1) GM pdfs and MMS models:* Let $p(X) = \sum_{z=1}^{M} c_z \cdot \mathcal{N}(\mu_z, \Sigma_z)$ be an $M$-term GM prior with component weights $c_z$ (such that $\sum_{z=1}^{M} c_z = 1$), component means $\mu_z$ and component covariances $\Sigma_z$, $\forall\ z \in \{1, .., M\}$. GMs are highly useful for dynamic estimation in multiple dimensions, as they can provide compact approximations to arbitrarily complex pdfs and enable closed-form predictions/updates in many systems with nonlinear/non-Gaussian dynamics and observation models [13].

As discussed in [14], the MMS model generalizes (2). MMS marginalizes over a set of softmax-distributed 'subclasses' in the discrete random variable $R$, whose realizations $r \in \{1, ..., S\}$ are each deterministically associated with a single class $j \in \{1, ..., m\}$ of $D$. Let $\sigma(j)$ be the set of $s_j$ subclasses belonging to class $D = j$. If $P(R = r|X)$ follows (2) and $P(D = j|R = r) = I(r \in \sigma(j))$ (i.e. the indicator function), then the law of total probability gives for $D = j$

$$P(D = j|X) = \sum_{r=1}^{S} P(R = r|X)P(D = j|R = r)$$

$$= \sum_{r \in \sigma(j)} P(R = r|X) = \frac{\sum_{r \in \sigma(j)} e^{w_q^T x + b_q}}{\sum_{g=1}^{S} e^{w_g^T x + b_g}} \quad (10)$$

Figure 2 shows the class probabilities (dashed lines) for an MMS model derived from the softmax model of Figure 1, where the class labels 'Next To', 'Nearby', and 'Far Away' in Figure 2 describe non-convex categories of a range-only observation (i.e. if the classes in Figure 1 represent subclasses for the classes of Figure 2, then $\sigma$('Next To') = {'Next To'}, $\sigma$('Nearby') = {'Near East','Near West'} and $\sigma$('Far From') = {'Far East','Far West'}). See [14] for details on learning MMS models from data.

*2) Bayesian data fusion:* To find the desired posterior $p(X|D = j)$, the discrete random variables $Z$ and $R$, which take values $z \in \{1, ..., M\}$ and $r \in \sigma(j)$, respectively, will be used to indicate realizations for latent GM components and MMS subclasses (where $P(Z = z) = c_z$ and $P(R = r|X)$ follows (2)). It can be shown that the joint pdf here is

$$p(D, X, Z, R) = P(D|R)P(R|X)p(X|Z)P(Z). \quad (11)$$

Using the law of total probability, the posterior is

$$p(X|D = j) = \sum_{Z=z, R=r} p(X|D = j, z, r)P(D = j, z, r) \quad (12)$$

For each pairing of $Z = z$ and $R = r$, consider the first summand term in (12); from Bayes' rule, (11) is

$$p(X|D = j, z, r) = \frac{P(D = j|r)P(r|X)p(X|z)P(z)}{\int P(D = j|r)P(r|X)p(X|z)P(z)dX}.$$

Canceling $P(z)$ and $P(D = j|r)$, we have

$$p(X|D = j, z, r) = \frac{P(r|X)p(X|z)}{\int P(r|X)p(X|z)dX} = \frac{P(r|X)p(X|z)}{P(r|z)}, \quad (13)$$

which is the posterior of $X$ conditioned on mixing component $z$ and subclass $r \in \sigma(j)$. The numerator of (13) is described by the product of a Gaussian $p(X|Z = z) = \mathcal{N}(\mu_z, \Sigma_z)$ and a softmax likelihood $P(R = r|X)$. Hence, for each prior component $z \in \{1, ..., M\}$ and MMS subclass observation $r \in \sigma(j)$, (13) can be well-approximated by a Gaussian $\mathcal{N}(\hat{\mu}_{zr}, \hat{\Sigma}_{zr})$ as in (4) using VBIS.

Next, consider the term $P(D = j, r, z)$ in (12), where for each fixed pairing of $z$ and $r$,

$$P(D = j, r, z) = \int p(X|z)P(r|X)P(D = j|r)P(z)dX,$$

$$= P(z) \int p(X|z)P(r|X)P(D = j|r)dX,$$

$$= P(z) \int p(X|z)P(r|X)dX, \quad (14)$$

where $P(z) = c_z$ and the last line follows from $P(D = j|r) = 1$ for $r \in \sigma(j)$, by definition of the MMS model. Note that $\int p(X|z)P(r|X)dX = \mathbb{E}\left[P(r|X)\right]_{p(X|z)} = P(r|z)$ is the expected softmax probability of subclass $r$ under component $z$; this can be directly approximated by

$$P(r|z) \approx \hat{P}(r|z) = \frac{1}{N_a} \sum_{r=1}^{N_a} P(R = r|X = x_r), \quad (15)$$

where $\{x_r\}_{r=1}^{N_a}$ is a set of $N_a$ samples taken directly from the GM prior component $\mathcal{N}(\mu_z, \Sigma_z)$. Hence, $P(D = j, r, z) \approx c_z \cdot \hat{P}(r|z)$, and (12) can be approximated as

$$\hat{p}(X|D = j) = \sum_{z,r} \beta_{zr} \cdot \mathcal{N}(\hat{\mu}_{zr}, \hat{\Sigma}_{zr}), \quad (16)$$

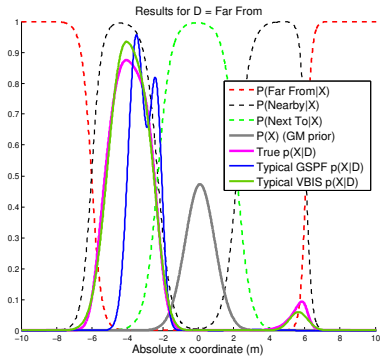


Fig. 2. 1D example for fusing a 'surprising' observation of $D$ = 'Far From' with MMS likelihood (dashed lines) and a GM prior (gray) with 4 components $[c_z, \mu_z, \Sigma_z]$: [0.9,0.1,0.7], [0.02,-0.1,0.7], [0.01,-0.7,1], [0.01,0.7,1]. The uncondensed 12-component posterior GM approximations using VBIS (green, with $N_s = N_a = 500$) and GSPF (blue, with $N_s = 1800$ to achieve the same computation time as VBIS) are plotted, showing that VBIS produces a better approximation to the true posterior (pink) than GSPF.

which is an $(M \cdot s_j)$-component GM with mixing weights $\beta_{zq} \propto c_z \cdot \hat{P}(r|z)$.

Eq. (16) is closely related to the output of a modified Gaussian sum particle filter (GSPF) [13], which performs GM-based state estimation using a parallel bank of particle filters. However, VBIS attempts to find an importance density $q$ for each term of (16) that is more efficient than the component prior $p(X|z)$, which the typical bootstrap GSPF (referred to hereafter as 'GSPF') uses instead to perform term-wise LWIS updates. When the prior and true posterior are similar (i.e. uninformative measurement updates), GSPF and VBIS fusion produce similar results (GSPF is somewhat faster since it does not require VB iterations). However, as shown in the 1D fusion example for a GM prior and MMS likelihood in Figure 2, GSPF is not robust to 'surprising' measurements. Note that, like the GSPF, the number of terms in (16) grow over time if $s_j > 1$ or if $M$ in $p(X)$ grows from dynamic state transitions. Standard GM compression methods (e.g. [15]) can counter this while minimizing information loss according to a suitable metric.

## IV. COOPERATIVE SEARCH APPLICATION

Motivated by the studies of [2], [3], [6], this section describes the application of the VBIS fusion method to a cooperative mission involving a real human-robot team searching for five hidden stationary targets. The goal is for a single human agent to assist a single autonomous mobile robot in correctly locating and identifying all targets as quickly as possible. The robot plans and navigates its own paths based on probabilistic target information, but has limited visual target detection capabilities. The human aids the robot by sharing new information (e.g. from visual or aural inspection of the search area) and confirming target detections, where the number of targets is known *a priori*.

### A. Experimental setup

Figure 3 (a) illustrates the physical setup for the target search over a 5 m x 10.5 m indoor area, featuring four 1 m long and four 2 m long obstacle walls. The walls

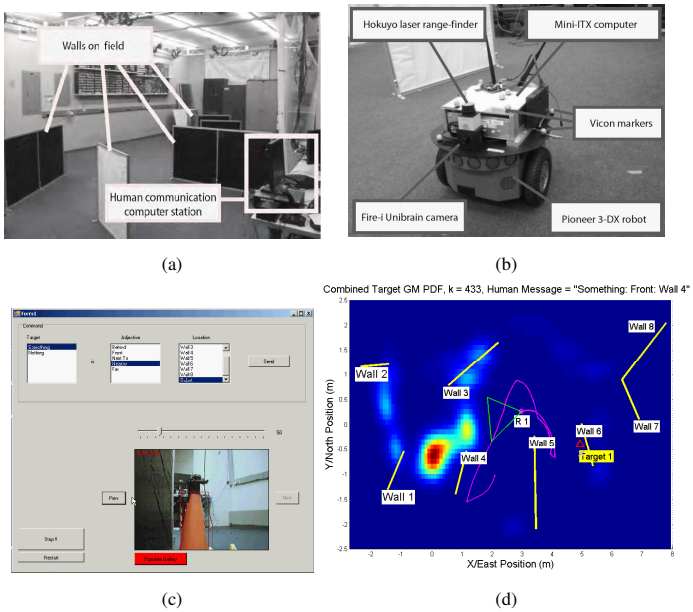


Fig. 3. (a) Field setup with walls and human computer station, (b) Pioneer 3-DX robot with sensing and computing hardware, (c) human-robot interaction GUI for sending human observations, (d) snapshot of combined GM pdf over field map during mission.

are placed such that the human (who remains seated off field at a computer) can only see a small portion of the search area by directly. The five targets were static orange traffic cones that were randomly placed behind walls. Each target location $X_u \in \mathbb{R}^2$, $u \in \{1, ..., 5\}$ is unknown to the human-robot team. As in [2], pdfs are used by the robot to autonomously plan search paths so that it can visually detect all targets without any direct human control inputs. Separate GM priors $p(X_u)$ are assumed for each $u$ at mission start, where $M = 15$ components are used and each $\mu_z^u \in \mathcal{U}([0, 5] \times [-1.5, 1.5])$ with $\Sigma_z^u = 3I$. The initial priors are highly uncertain and possibly inconsistent with the true locations of some targets, so the robot updates the pdfs by fusing two information sources: (1) $\mathcal{V}^k$, the set of all binary readings from an onboard visual target detector up to time $k$, and (2) $\mathcal{H}^k$, the set of all categorical target location observations provided by the human up to time $k$. Section IV-B details the information fusion process. As this work focuses on the data fusion problem, a simple sub-optimal greedy search strategy is used with the updated pdfs.

Figure 3 (b) shows the Pioneer 3-DX autonomous mobile robot used in the study. It uses Vicon motion-capture markers for accurate pose estimation, a laser range finder for obstacle avoidance and target localization, and an onboard camera with software that detects cones up to 1 m with a 42.5 deg field of view. The robot moves at a constant speed of 0.3 m/s and uses a $D^*$ algorithm to autonomously plan search paths based on the updated combined target GM pdf,

$$p(X_{comb,k}) = \sum_{u \in U_k} \frac{1}{|U_k|} \cdot p(X_u | \mathcal{V}^k, \mathcal{H}^k), \qquad (17)$$

where $U_k$ is the set of undetected targets at time $k$. The robot discretizes (17) and selects the non-obstacle grid cell with the highest pdf value as its next goal point. The robot

automatically replans and follows a new path if it either reaches a goal point or a receives a new human observation. The robot also streams camera images to the human at 1 Hz and pauses to report target detections, which are verified by the human. False alarms are logged to prevent reacquisition, while true detections result in the removal of the GM for target $u$ from (17). The human communicates with the robot via the GUI shown in Fig. 3 (c) and can also see an updated 2D surface plot of (17) during the mission (Fig. 3 (d)). The human and robot share a common fixed map of the search area to give consistent contextual information for fusion.

### B. Online fusion updates for GM pdfs

From Bayes' rule, the target $u$ pdf given new vision observations $C_k$ and human observations $D_k$ at time $k$ is

$$p(X_u|\mathcal{V}^k, \mathcal{H}^k) \propto p(X_u|\mathcal{V}^{k-1}, \mathcal{H}^{k-1}) P(C_k|X_u) P(D_k|X_u),$$

where $p(X_u|\mathcal{V}^{k-1}, \mathcal{H}^{k-1})$ is the static target posterior from step $k-1$, which equals $p(X_u)$ at $k=1$. Bayes' rule also divides this into sequential fusion updates for $C_k$ and $D_k$,

$$p(X_u|\mathcal{V}^k, \mathcal{H}^{k-1}) \propto p(X_u|\mathcal{V}^{k-1}, \mathcal{H}^{k-1}) P(C_k|X_u), \quad (18)$$

$$p(X_u|\mathcal{V}^k, \mathcal{H}^k) \propto p(X_u|\mathcal{V}^k, \mathcal{H}^{k-1}) P(D_k|X_u). \quad (19)$$

Eq. (18) is skipped for false detections, which are assumed to be filtered out perfectly by the human. Eq. (19) is skipped at time $k$ if $D_k$ is empty. $P(C_k|X_u)$ is given by the 2D MMS model shown in the top row of Figure 4 (a), which models visual detection probability as a function of robot position and camera field of view. Since the robot's slow motion means that prior and posterior GM pdfs are similar between frequent 'no detection' readings, the update in (18) runs at 2 Hz using GSPF with 1000 samples per component update (this gives nearly identical results to the somewhat slower VBIS method).

Each human observation $D_k$ is selected from a collection of mutually exclusive categorical descriptions of $X_u$ with MMS likelihoods $P(D_k|X_u)$. Human observations are sent using the 3-field message structure '(*Something/Nothing*) is (*Preposition*) (*Location*)', where field entries are set via menus in the GUI. The first field allows the human to provide either positive or negative soft information. The second field determines whether an observation is relative to one of the walls in the search area or the robot's current location. The final field entry comes from either a set of relative ranges {'Next To', 'Nearby', 'Far'} or a set of wall-relative location descriptions {'Front Of', 'Behind', 'Inside'}. Each set is described by a separate MMS model (learned offline) whose origin can be shifted/rotated to match the *Location* coordinates. The bottom row of Figure 4 (b) shows a typical MMS model for 'Something Nearby the Robot'. Since inconsistent information in $p(X_u)$ can make $D_k$ seem 'surprising', (19) is performed with VBIS, where ambiguities from the use of 'Something'/'Nothing' descriptors are handled via naive probabilistic data association. Following updates via (18) and (19), each target GM is condensed to 15 components via Salmond's method [15], which preserves the overall GM mean and covariance.

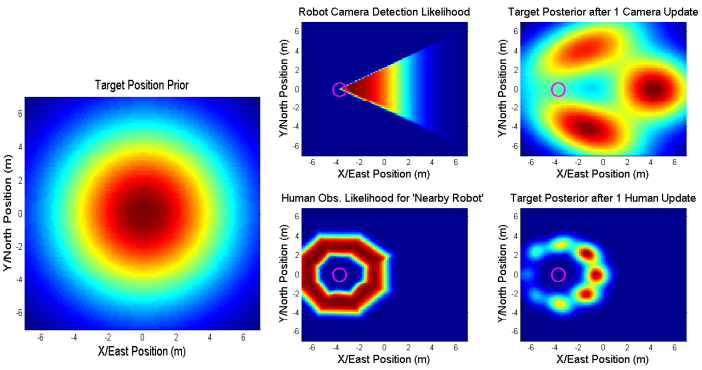


Fig. 4. (a) Example Gaussian target location prior, (b) *Top row*: MMS camera detection model, and 3 component posterior GM pdf from GSPF after fusing 'no detection' report from vision sensor with prior in (a). *Bottom row*: MMS human observation model for D='Target is Nearby robot', and 8 component posterior GM pdf from VBIS after fusion with prior in (a).

### C. Results

Three search missions of 20 minute nominal length were conducted with different target locations and priors. In the first two trials, the human-robot team found all 5 targets in under 15 minutes. In the third trial, 4 targets were found in 15 minutes and the remaining target was found after 21 minutes due to difficulties with obstacle navigation. The best baseline greedy search under the same prior conditions without human information fusion (but with human target confirmation) found 2 targets in 15 minutes and 3 targets after 20 minutes (an added 3 minutes yielded no new detections). An average of 61 human messages were sent per trial, with 74 messages sent in the last trial. This reflects the fact that the impact of new positive human information was downweighted by probabilistic data association, so that the human had to adjust inconsistent prior information about target locations by repeatedly sending positive messages to 'convince' the robot about new information. For instance, Figure 5 shows snapshots of (17) from the first trial as a sequence of human messages shifted the GM peaks to the back of Wall 2, where a target was successfully found by the robot (this target was undetected in the baseline trial without human fusion, as the pdf behind Wall 2 remained very small). Hence, even with imperfect priors, a naive greedy search strategy, limited soft observations in $\mathcal{V}^k$ and $\mathcal{H}^k$, soft human information fusion was quite beneficial.

To assess the accuracy of the VBIS for fusing $\mathcal{H}^k$ alongside GSPF fusion of $\mathcal{V}^k$, the logged fusion pdfs were compared to offline pdfs obtained by using GSPF to fuse both $\mathcal{H}^k$ and $\mathcal{V}^k$. Using an offline grid-based estimate with 0.1 m cell resolution as the 'ground truth' fusion pdf at each time step, the Kullback-Leibler divergences (KLDs) for the VBIS and GSPF pdfs with respect to the truth were computed at each time step for the first trial. The KLD between two pdfs $g(X)$ and $h(X)$ at any time step is

$$\mathrm{KL}(g||h) = \int_X g(X) \log \frac{g(X)}{h(X)} dX, \quad (20)$$

where $g(X)$ is the true posterior. The KLD can be interpreted

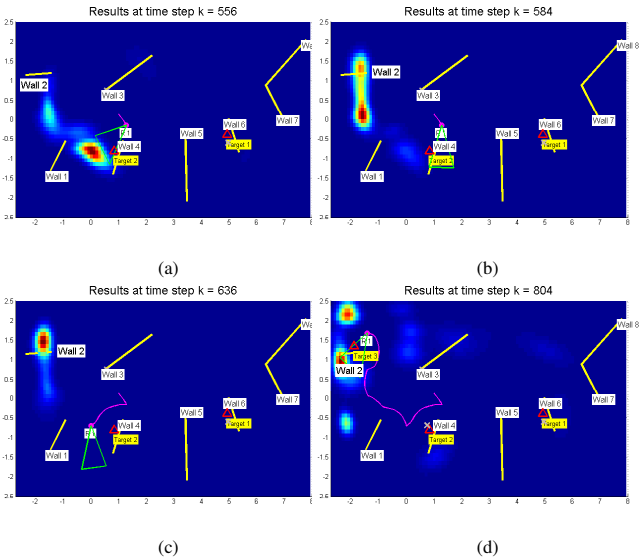


Fig. 5. Combined target pdf shifting behind Wall 2 for sequence of human message updates: (a) pdf at k = 556 secs, (b) after 7 human messages, (c) after 4 more messages, (d) after target 3 found behind Wall 2.

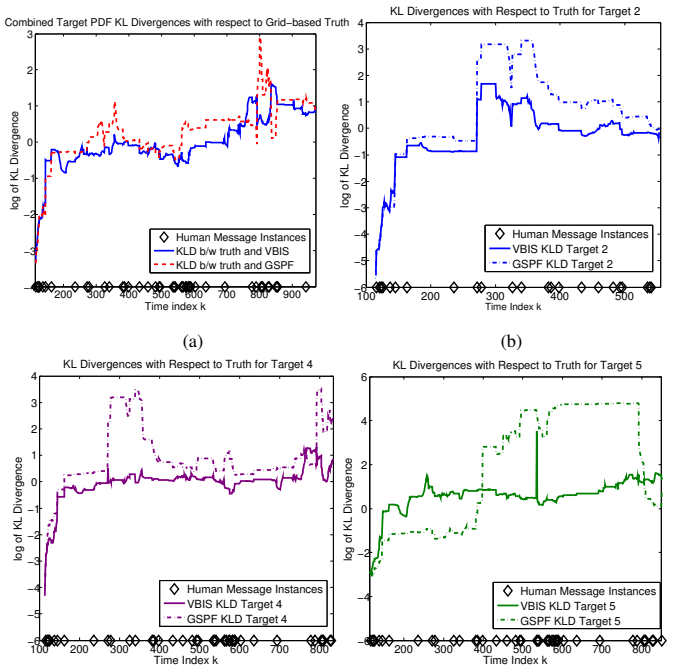


Fig. 6. (a) Log KLDs with respect to grid-based true posterior pdf for VBIS (solid) and GSPF (dashed) GM estimates of total target search distribution vs. time. (b)-(d) Log KLDs for targets 2, 4 and 5 vs. time.

as a 'distance measure' between two pdfs, where $h$ is a good approximation to $g$ if (20) is near zero. The KLDs for (17) over time are shown in Figure 6 (a) (log values shown). It can be seen that the results using VBIS for $\mathcal{H}^k$ fusion are generally closer to the truth than the GSPF fusion results. Note that the use of finite GMs to approximate $g(X)$ is expected to produce non-zero KLDs, since GMs can add extra probability mass back into tails where $g(X)$ is actually much smaller (especially after Salmond merging). Figures 6 (b)-(d) show individual target pdf KLDs for targets 2, 4, and 5 (log values shown), which are often much larger with GSPF-only fusion than with VBIS. Closer inspection revealed that the GSPF-only fusion pdfs sometimes dropped GM components that corresponded to significant modes of the true fusion pdfs and is thus heavily penalized in (20). On the other hand, VBIS fusion retains all modes of the true posterior following fusion of $\mathcal{H}^k$, although some regions of $X_u$ are underweighted relative to truth (likely as the result of Salmond merging). Note that since the logged data was collected in closed-loop with VBIS for fusing $\mathcal{H}^k$, the GSPF-only pdfs sometimes recover from mode losses accidentally, causing severe fluctuations in the corresponding KLDs.

## V. CONCLUSIONS

The VBIS method was presented for hybrid Bayesian fusion of soft categorical observations of continuous states. VBIS was developed for the baseline case of Gaussian priors with softmax likelihoods and then extended to the more general case of Gaussian mixture priors with multimodal softmax likelihoods. An experimental multitarget search mission involving a real human-robot team demonstrated the reliability and usefulness of the proposed fusion approach. Future work will study the extension to decentralized GM fusion for more complex human-robot team architectures.

## REFERENCES


[1] P. Bladon, P. Day, T. Hughes, and P. Stanley, "High-level fusion using Bayesian networks: Applications in command and control," in *Information Fusion for Command Support*, 2004, pp. p4.4–p4.18.

[2] F. Bourgault, "Decentralized control in a Bayesian world," Ph.D. dissertation, University of Sydney, 2005.

[3] T. Kaupp, B. Douillard, F. Ramos, A. Makarenko, and B. Upcroft, "Shared environment representation for a human-robot team performing information fusion," *Journal of Field Robotics*, vol. 24, no. 11, pp. 911–942, 2007.

[4] W. Koch, "On 'negative' information in tracking and sensor data fusion: Discussion of selected examples," in *FUSION 2004*, pp. 91–98.

[5] J. Hoffmann, M. Spranger, D. Gohring, and M. Jungel, "Negative information and proprioception in monte carlo self-localization for 4-legged robots," in *19th Int'l Joint Conf.on AI (IJCAI)*, 2005.

[6] M. Lewis, H. Wang, P. Velgapudi, P. Scerri, and K. Sycara, "Using humans as sensors in robotic search," in *FUSION 2009*, pp. 1249–1256.

[7] F. Bourgault, A. Chokshi, J. Wang, D. Shah, J. Schoenberg, R. Iyer, F. Cedano, and M. Campbell, "Scalable Bayesian human-robot cooperation in mobile sensor networks," in *IROS 2008*, pp. 2342–2349.

[8] N. Ahmed and M. Campbell, "Variational Bayesian data fusion of multi-class discrete observations with applications to cooperative human-robot estimation," in *ICRA 2010*, pp. 186–191.

[9] C. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer, 2006.

[10] Y. Bar-Shalom, X. Li, and T. Kirubarajan, *Estimation with Applications to Navigation and Tracking*. New York: Wiley, 2001.

[11] G. Bouchard, "Efficient bounds for the softmax function and applications to approximate inference in hybrid models," in *NIPS 2007 Workshop for Approximate Bayesian Inference in Continuous/Hybrid Systems*, Whistler, BC, Canada, 2007.

[12] M. Arulumpalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, 2002.

[13] J. Kotecha and P. Djuric, "Gaussian sum particle filtering," *IEEE Trans. on Sig. Proc.*, vol. 51, no. 10, pp. 2602–2612, 2003.

[14] N. Ahmed and M. Campbell, "Multimodal operator decision models," in *American Control Conference 2008*, 2008, pp. 4504 – 4509.

[15] D. Salmond, "Mixture reduction algorithms for uncertain tracking," Farnborough, UK: Royal Aerospace Est., Tech. Rep. 88004, 1988.