2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

Variational Bayesian Data Fusion of Multi-class Discrete Observations with Applications to Cooperative Human-Robot Estimation

Nisar Ahmed and Mark Campbell
Sibley School of Mechanical and Aerospace Engineering
Cornell University, Ithaca, NY 14853, USA
{nra6,mc288}@cornell.edu

***Abstract*— A new method is presented for fusing conventional continuous sensor observations with discrete multi-categorical state-dependent information, which can be furnished by humans in many cooperative human-robot interaction problems. The hybrid likelihood function for mapping between continuous hidden states and categorical observations are specified via *softmax* models. Although softmax models avoid discretization of continuous states, they are challenging to implement for real-time data fusion since they are not analytically integrable. An approximation based on variational Bayesian (VB) methods is presented here to obtain fast closed-form Gaussian solutions to the desired posteriors in cases where the hidden continuous states have Gaussian pdfs. A joint human-robot target localization example illustrates the properties and utility of the VB hybrid fusion strategy, which also applies more generally to inference in hybrid Bayesian networks and mixture models.**

I. Introduction

Problems of hybrid (i.e. continuous and discrete) estimation are often encountered in human-robot interaction (HRI) applications; since humans are bound to discretely categorize the state of the world as an efficient way to process and convey information, proper characterization and use of humans as hybrid information sources is an essential step toward improving HRI and HR team performance [1], [2]. With cooperative HR applications such as target tracking and search-and-rescue in mind, refs. [3], [4] and [5] explored direct incorporation of raw human sensory observations for improving dynamic recursive Bayesian state estimation (RBE). However, the 'human sensor models' used in these studies cannot describe the generally nonlinear/non-Gaussian relationships between continuous states and non-binary categorical human observations.

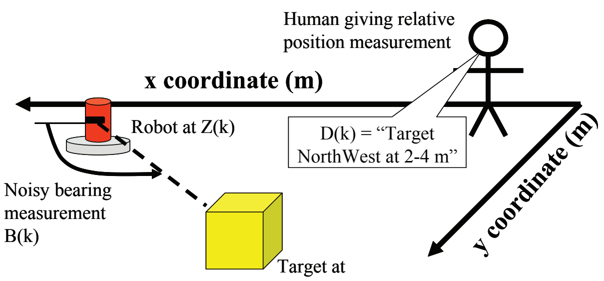

Fig. 1. Cooperative 2D localization of a static target by a human-robot team. The index k denotes discrete time.

Motivated by [1], [3]-[5], Figure 1 illustrates a cooperative search problem involving 2D localization of a static target by a mobile robot and a human agent. The robot moves through the search space and takes continuous bearings-only measurements of the target's position, while the human independently reports the approximate relative range and bearing between the robot and target using visual inspection. Clearly, the human's observations are not expected to be as precise as, say, laser range measurements. Rather, the human is expected to report observations that discretely categorize the relative bearing (e.g. into canonical directions such as 'North', 'NorthEast', 'East', etc.) and the relative range (e.g. in 2 m bins up to some maximum distance). For example, the human might report that the target is '2-4 m NorthWest' of the robot's current location, or (in more natural terms) that the target is 'nearby' the robot and 'in front, to the left.'

The human's observations can significantly improve the target location estimate, which can be very uncertain if found solely via the robot's bearings sensor [6]. Therefore, it is desired to *fuse* prior information of the target position with the robot's continuous measurements and the human's discrete observations via Bayes' rule to update the target position belief. From Figure 1, suppose for now that $p(X|B,Z)$ represents the robot's belief of the target location given its own available position and bearing measurements. If $p(D|X,Z)$ is the human's discrete observation likelihood given the true target and robot locations, Bayes' rule gives

$$p(X|D,B,Z) = \frac{p(D|X,Z)\ p(X|B,Z)}{\int p(D|X,Z)\ p(X|B,Z)\ dX} \quad (1)$$

This raises two questions: (i) how to specify $p(D|X,Z)$, and (ii) how to compute (1) for efficient RBE? Answering (i) requires finding a function that maps continuous robot and target states to discrete probability mass functions. The *softmax function* is such a distribution that is popular for hybrid probabilistic models over m discrete categories [7],

$$p(D=j|Q) = \frac{e^{w_j^T q + b_j}}{\sum_{k=1}^{m} e^{w_k^T q + b_k}}, \quad (2)$$

where $Q = [Z,X]^T$, w_k is the vector weight for class k and b_k is a scalar bias, for $j,k \in \{1,...,m\}$. However, the normalizing integral in (1) becomes

$$\int \frac{e^{w_j^T x + b_j}}{\sum_{k=1}^{m} e^{w_k^T x + b_k}} \; p(X|B,Z) \; dX, \tag{3}$$

which unfortunately cannot be evaluated analytically for any $p(X|B,Z)$. Therefore, (1) has no *exact* closed-form solution, which prevents efficient *exact* hybrid RBE updates for X (e.g. as in Kalman filtering for linear-Gaussian systems). However, it is possible to obtain good *approximations* for this hybrid Bayesian inference problem.

This paper presents a new approximate hybrid inference approach based on variational Bayesian (VB) methods [8], [9], [10]. The method applies to a broad class of system models that use the softmax function to model discrete sensor observations of states with Gaussian priors. The proposed VB approximation yields fast, closed-form Gaussian pdf posterior approximations, thus enabling efficient hybrid data fusion for RBE. This VB method is also generally useful for efficient inference in hybrid Bayesian networks [7], [8], which have become increasingly useful for HRI [3], [11], [12], [13]. Section II gives more background and Section III details the VB method and its properties; Section IV demonstrates these via the HR cooperative localization problem. Section V gives conclusions and future work.

II. Hybrid Probabilistic Models and Data Fusion

A. Softmax Likelihoods for State-dependent Discrete Sensors

The softmax model is often used in statistical pattern classification [9], where it is desired find boundaries between discrete classes as a function of an input space. The log-odds of (2) gives the probabilistic 'class boundaries' as linear hyperplanes in Q, where the weights control the location of these boundaries for fixed Q and the biases shift the boundaries from the origin. Softmax weights and biases can be easily estimated from training data; in addition, mixture-based extensions of (2) can be used to approximate more complex nonlinear class boundaries (e.g. see [14]). For clarity, we focus here on 'basic' softmax models. As this paper focuses on hybrid fusion methods, we refer the reader to [7], [14] for more details on hybrid probabilistic modeling.

B. RBE and Hybrid Bayesian Data Fusion

Using (1) as an example, if B and D could always be *approximately* fused with a Gaussian prior for X to yield a Gaussian posterior, then Gaussian sufficient statistics need only be maintained for efficient RBE of X [6]. This allows for the fact that good Gaussian approximations still exist even when exact fusion for B alone is impossible (e.g. EKFs [6]). To retain Gaussian posteriors for hybrid data fusion, Monte Carlo techniques can be used [15]. However, these require computationally intensive methods to ensure accuracy, which can hamper real-time execution. While discretization is a 'simple' solution, it scales poorly and loses the Gaussian representation. Ref. [8] gives a VB approximation to the *binary* discrete inference/fusion problem in hybrid Bayesian networks. This VB method approximates the true joint pdf as the product of the prior and a Gaussian lower bound

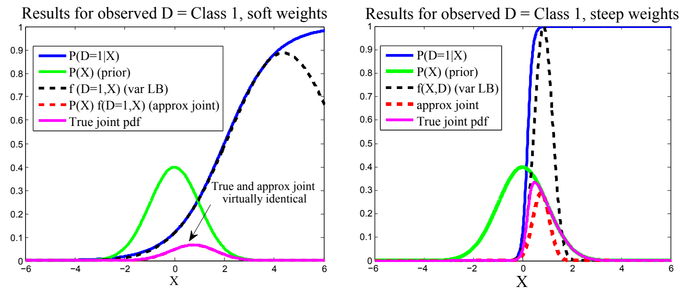

(a) (b)

Fig. 2. For standard normal prior (green): (a) likelihood (blue) is well-approximated by Gaussian bound (black), so approximate joint pdf is very close to true one (magenta=blue×green); (b) likelihood approximation degrades at steeper softmax weights, but joint pdf is still well-approximated.

to a binary hybrid likelihood. While this method gives fast closed-form solutions that are good Gaussian approximations to the true posterior, it is limited to $m = 2$. We present a generalization of the VB method for $m \geq 2$ next.

III. VB Fusion for Multi-class Observations

A. General hybrid inference problem formulation

Let the unobserved random state vector $X \in \mathbb{R}^n$ have Gaussian prior $p(X) = \mathcal{N}(\mu, \Sigma)$ with known mean $\mu \in \mathbb{R}^n$ and covariance $\Sigma \in \mathbb{R}^{n \times n}$. Let D be an m-valued discrete sensor variable for some integer $m \geq 2$, where $p(D = j|X)$ follows (2) with known weights $w_j \in \mathbb{R}^n$ and biases $b_j \in \mathbb{R}$ for $j \in \{1, ..., m\}$. We assume here that all available continuous sensor measurement updates have already been performed via Bayes' rule, so that $p(X)$ is generally a Gaussian posterior from purely continuous fusion. The posterior pdf of X given $D = j$ is thus

$$p(X|D=j) = \frac{1}{C} \; p(X, D=j) = \frac{1}{C} \; p(X) \; p(D=j|X)$$
$$= \frac{1}{C} \left|2\pi\Sigma^{-1}\right| e^{-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)} \frac{e^{w_j^T x + b_j}}{\sum_{c=1}^{m} e^{w_c^T x + b_c}}, \tag{4}$$

where $C \equiv \int p(X) p(D=j|X) dX$. Inspection of (4) reveals that the joint pdf is neither Gaussian nor analytically integrable with respect to X, due to the denominator of (2).

B. Derivation of Approximate Posterior

We now derive a closed-form variational Gaussian approximation to $p(X|D=j)$ based on the idea that the joint pdf $p(X, D=j) = p(X) \; p(D=j|X)$ is well-approximated by an *unnormalized* Gaussian function over X; this is illustrated for the binary softmax (i.e. logistic) case [8] in Figure 2.

Replacing the softmax function (2) with an unnormalized Gaussian $f(D=j, X)$ gives

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

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

Note that now $\hat{p}(X, D=j)$ is an unnormalized Gaussian and $\hat{C}$ is closed-form. This leads to a closed-form normalized Gaussian approximation to the posterior,

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

where $\hat{\mu}$, $\hat{\Sigma}$ and $\hat{C}$ can be obtained by inspection of (5). For $m \geq 2$, $f(D=j,X)$ can be obtained via the variational softmax lower-bound proposed in [10]. This bound is obtained by replacing the denominator in (2) with a variational upper bound composed of a product of m unnormalized Gaussians. Namely, for any set of scalars α, ξ_c and y_c for $c \in \{1,...,m\}$, [10] shows that

$$\log\left(\sum_{c=1}^{m} e^{y_c}\right) \leq \alpha + \sum_{c=1}^{m} \frac{y_c - \alpha - \xi_c}{2} + \lambda(\xi_c)[(y_c-\alpha)^2 - \xi_c^2] + \log(1+e^{\xi_c}) \quad (7)$$

where $\lambda(\xi_c) = \frac{1}{2\xi_c}\left[\frac{1}{1+e^{-\xi_c}} - \frac{1}{2}\right]$. The variables α and ξ_c are free *variational parameters* for minimizing the upper bound in (7) for known y_c to provide the tightest possible (lower bounding) approximation to the original softmax function (2). Zeroing the derivatives with respect to α and ξ_c gives

$$\xi_c^2 = y_c^2 + \alpha^2 - 2\alpha y_c, \quad \alpha = \frac{(\frac{m-2}{4}) + \sum_{c=1}^{m}\lambda(\xi_c)y_c}{\sum_{c=1}^{m}\lambda(\xi_c)}. \quad (8)$$

Here, $y_c = w_c^T x + b_c$; however, since X is unobserved, each y_c is an unknown random variable. To handle this, an efficient procedure for finding α and ξ_c that maximize the 'average' value of (7) is presented in the next section. Assuming for now that α and all ξ_c are fixed, (7) can be used to obtain an approximate Gaussian posterior. From (2),

$$\log p(D=j|X) = w_j^T x + b_j - \log\left(\sum_{c=1}^{m} e^{w_c^T x + b_c}\right). \quad (9)$$

Replacing the last term with (7) (with $y_c = w_c^T x + b_c$), simplifying terms and exponentiating, we get

$$f(D=j,x) = \exp(g_j + h_j^T x - \frac{1}{2}x^T K_j x), \quad (10)$$

$$\text{where} \quad g_j = \frac{1}{2}\left[b_j - \sum_{c\neq j} b_c\right] + \alpha(\frac{m}{2} - 1)$$
$$+ \sum_{c=1}^{m} \frac{\xi_c}{2} + \lambda(\xi_c)\left[\xi_c^2 - (b_c-\alpha)^2\right] - \log(1+e^{\xi_c}), \quad (11)$$

$$h_j = \frac{1}{2}\left[w_j - \sum_{c\neq j} w_c\right] + 2\sum_{c=1}^{m}\lambda(\xi_c)(\alpha - b_c)w_c, \quad (12)$$

$$K_j = 2\sum_{c=1}^{m}\lambda(\xi_c)w_c w_c^T, \quad (13)$$

and $f(D=j,X) \leq p(D=j|X)$ from (7). Note that $f(D=j,X)$ is an *unnormalized* Gaussian over X. Now,

$$p(X) = \exp(g_p + h_p^T x - \frac{1}{2}x^T K_p x), \quad (14)$$

where $g_p = -\frac{1}{2}(\log[|2\pi\Sigma|] + \mu^T K_p \mu)$, $h_p = K_p\mu$ and $K_p = \Sigma^{-1}$. Substituting (10) and (14) into (5) therefore gives

$$p(X,D=j) \geq \hat{p}(X,D=j) = p(X)\ f(D=j,X) \quad (15)$$
$$= \exp(g_l + h_l^T x - \frac{1}{2}x^T K_l x) \quad (16)$$

where $g_l = g_p + g_j$, $h_l = h_p + h_j$, and $K_l = K_p + K_j$. This an unnormalized Gaussian approximation to the joint pdf; normalization gives the desired approximation to the posterior $\hat{p}(X|D=j) = \mathcal{N}(\hat{\mu},\hat{\Sigma})$, where

$$\hat{\Sigma} = K_l^{-1}, \quad \hat{\mu} = K_l^{-1}h_l. \quad (17)$$

Interestingly, these updates to the Gaussian state prior are very similar to the updates of a state-space information filter for linear-Gaussian systems [6]. In particular, the information matrix update in (17) is independent of the observed sensor value. Also, since h_l is analogous to an information state, the softmax weights w_j can be viewed as the 'average information' contained within each class j about X. Indeed, the degree of ambiguity between the discrete classes given a particular state $X = x$ is analogous to continuous sensor noise. Larger weights imply more information and less ambiguity for a given x (i.e. less noise between discrete classes), while softer weights imply the opposite. If class j is 'well-separated' from all other classes throughout the state-space X, then class j observations provide much information about X. If observed class j is not easily distinguished from other classes, j contributes little to (17) via (12) and (13). Finally, it should be noted that this approximation is 'variational' since it can be shown that $\hat{p}(X|D=j)$ minimizes the Kullback-Leibler divergence (KLD) functional $\text{KL}[G(X)||p(X|D=j)]$ between any Gaussian pdf $G(X)$ derived via (7) and the true (non-Gaussian) posterior.

C. Variational Parameter Optimization

To find α and ξ_c, we maximize the *approximate* marginal log-likelihood of the observed data (i.e. the evidence),

$$\log\hat{C} = \log\int p(X)\ f(D=j,X)\ dX, \quad (18)$$

where $\log\hat{C} \leq \log C$. Since (18) resembles a marginal data log-likelihood, the EM algorithm can be invoked to efficiently find the optimal α and ξ_c along with $\hat{p}(X|D=j)$. In the E-step, $\hat{p}(X|D=j)$ is obtained from (17) for fixed α and ξ_c; in the M-step, $\hat{p}(X|D=j)$ is held fixed and used to re-estimate α and ξ_c from (8) using the *expected values* of y_c. In this case, EM always leads to a unique VB posterior $\hat{p}(X|D=j)$; while proof is omitted due to limited space, this stems from two facts: (i) VB approximations monotonically converge to local maxima of the true posterior [9] and (ii) $\log p(X|D=j)$ is convex. To gauge the EM algorithm's monotonic convergence, (18) can be computed after each M-step as

$$\log\hat{C} = \langle y_j\rangle - \alpha + \sum_{c=1}^{m}\left\{\frac{1}{2}(\alpha + \xi_c - \langle y_c\rangle)\right.$$
$$\left. -\lambda(\xi_c)[\langle y_c^2\rangle - 2\alpha\langle y_c\rangle + \alpha^2 - \xi_c^2] - \log(1+e^{\xi_c})\right\}$$
$$-\frac{1}{2}\left(\log\frac{|\Sigma|}{|\hat{\Sigma}|} + \text{tr}(\Sigma^{-1}\hat{\Sigma}) + (\mu-\hat{\mu})^T\Sigma^{-1}(\mu-\hat{\mu}) - n\right),$$

where $\langle y_c\rangle = w_c^T\hat{\mu} + b_c$,
$$\langle y_c^2\rangle = w_c^T\left(\hat{\Sigma} + \hat{\mu}\hat{\mu}^T\right)w_c + 2w_c^T\hat{\mu}b_c + b_c^2$$

TABLE I

VB HYBRID FUSION EM ALGORITHM

0. *Given*: μ, Σ, $D=j$, known softmax weights w_k and biases b_k, arbitrary α and arbitrary ξ_c for $j,c,k \in \{1,...,m\}$
1. **E-step**: for all fixed ξ_c and α,
(a) compute $\hat{\mu}$ and $\hat{\Sigma}$ (eq. 17)
(b) compute $\langle y_c \rangle$ and $\langle y_c^2 \rangle$
2. **M-step**: for all fixed $\langle y_c \rangle$ and $\langle y_c^2 \rangle$,
for i=1:n_{lc}
(a) compute all ξ_c for fixed α (eq. 8)
(b) compute α for all fixed ξ_c (eq. 8)
end
3. Compute $\log \hat{C}$; if not converged, Repeat 1 and 2. Otherwise, stop.

Table I summarizes the VB hybrid fusion EM algorithm. Note that α and ξ_c are non-linearly coupled in the M-step via (8); this simply requires an extra inner-loop for iterative convergence, where typically $n_{lc} \leq$15. The number of EM steps needed for convergence (defined by some tolerance on the change in $\log \hat{C}$) will vary with the application, though the computations can be done in real-time, as shown later.

D. Properties

The VB hybrid inference method for *binary* softmax likelihoods ($m=2$) has several desirable properties [8],[16]: (i) its closed-form Gaussian approximation is fast and easy to compute, (ii) its posterior mean is very close to the true posterior mean, (iii) it is insensitive to the observation likelihood, and (iv) it is unique (i.e. no local maxima).

These properties also hold for the proposed softmax VB method in the general case of $m \geq 2$. Property (i) implies that real-time approximate hybrid data fusion is feasible without costly discretization or numerical integration of the state variables. Since the Gaussian approximate posterior under hybrid observation updates ensures that all future hybrid updates yield (approximate) Gaussians, only the sufficient statistics need to be maintained for efficient RBE.

Property (ii) implies that the VB approximate mean is an excellent approximate MMSE estimator with hybrid observations, since the optimal MMSE estimator is the true posterior mean [6]. Note that this gives VB methods an advantage over MAP-based Gaussian posterior approximations, which can place too much probability mass in regions of low support near asymmetric modes (e.g. see Figure 2(b)). VB Gaussians 'compensate' for posterior asymmetry by shifting closer to the true posterior mean [16].

Property (iii) implies robustness to unlikely evidence which 'disagrees' with the prior or is 'surprising' given its low probability. This gives VB a key advantage over fast importance sampling methods, which degenerate when samples are drawn far from the true posterior [15]. Markov Chain Monte Carlo sampling methods can overcome this, but can be too slow for real-time use. Finally, property (iv) guarantees that the EM algorithm in Table I converges to a unique $\hat{p}(X|D=j)$ for any initial choice of α and ξ_c. This is important for cases in which complex priors are modeled as Gaussian mixtures and complex likelihoods are mixtures of softmax models, since all modal contributions in the resulting posterior can be tracked. Posterior mode loss often occurs with Monte Carlo approximations and has important consequences in practice (e.g. see [17]).

IV. APPLICATION TO HR COOPERATIVE ESTIMATION

A. Model Formulation

Figure 3 shows the hybrid Dynamic Bayesian Network (DBN) model of the HR cooperative localization problem from Figure 1, where X is the 2D position of the static target. It is assumed that the robot's position $Z(k)$ is always known at time step k, so that $p(Z(k)|Z(k-1))$ is arbitrary (the case of unknown $Z(k)$ is considered later). The robot bearing measurement $B(k)$ has a known likelihood $p(B(k)|Z(k),X)$ that admits a Gaussian 'partial posterior' $p(X|Z_{1:k}, B_{1:k}, D_{1:k-1})$. The human's sensor model likelihood $p(D(k)|Z(k),X)$ is given by (2) for some discrete label set $\mathcal{D}$; this model is assumed without loss of generality to be independent of the human's location. We now illustrate the VB hybrid fusion method via different localization scenarios with $m>2$. For clarity, we use different Gaussian target priors and softmax models in each case, and only focus on the discrete sensor update at a single time step. For comparison, we also examine fusions obtained by 'exact' methods and multiple trials of likelihood weighting (LW) importance sampling [7]. Note that LW behaves very similarly to bootstrap particle filtering, since the measurement likelihood is used as the importance weight [15].

B. 1D Localization Scenario

Assume the y coordinate of X is known to be $y=0$ and that $Z(1) = [0,0]^T$; the discrete fusion examined here occurs at $k=1$ *before* fusion of any continuous $B(k)$. The human provides $D=j$, where $j \in \mathcal{D} =$ {'Far West', 'Near West', 'Next To', 'Near East', 'Far East'} (m =5). From the DBN model and Bayes' rule,

$$p(X|D=j,Z_1) = \frac{p(X)p(D=j|X,Z_1)}{\int p(X)p(D=j|X,Z_1)dX}. \quad (19)$$

Figs. 4 (a)-(c) show the true and approximate posteriors for each of the fusion methods for three different scenarios, each with a different discrete observation j and Gaussian prior (the true X in each case is not shown for clarity of illustration). The softmax likelihoods for each class are shown as dashed lines. The observations j in (a)-(c) are 'Near West', 'Next To', and 'Near East', respectively; the true posteriors (magenta) were obtained by numerical integration.

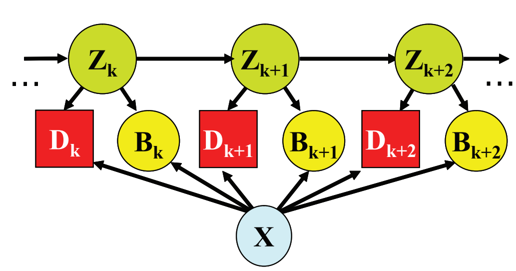

Fig. 3. Hybrid DBN model of the human-robot localization problem. The joint pdf is given by the product of the local conditional distributions for each node. Only the static target location node X is unobserved.

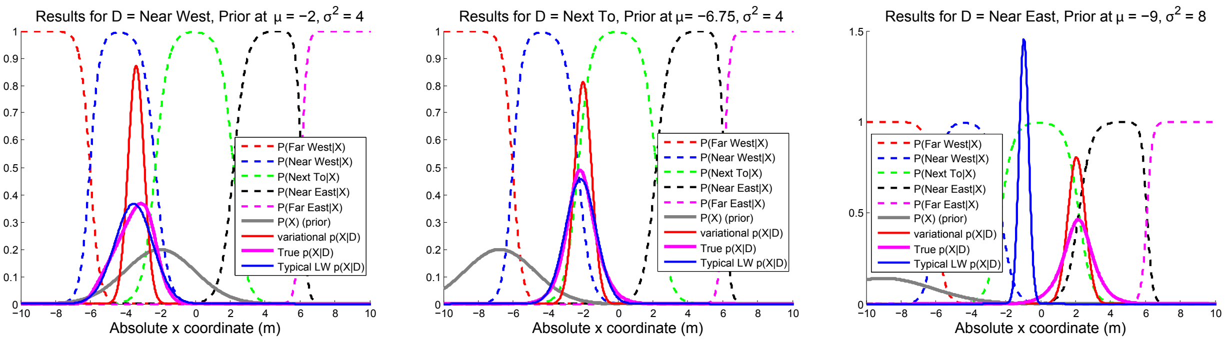

(a) (b) (c)

Fig. 4. Illustration for 1D cases.

Case	Prior μ,σ^2	True Posterior μ,σ^2	VB posterior μ,σ^2	LW posterior μ,σ^2	VB time, EM steps	LW time, SE
a	-2,4	-3.5155, 1.1386	-3.3862, 0.2097	-3.5458±0.0640, 1.0913±0.0252	1.6e-02 s, 18	1e-03 s, 50 %
b	-6.75,4	-2.0717,0.8594	-1.9949, 0.2412	-2.1571+/-0.2041, 0.8969+/-0.1314	1.6e-02s, 16	1e-03 s, 15 %
c	-9, 8	2.1404, 1.0087	2.0559, 0.2443	-0.4659+/-0.8105, 0.2488+/-0.0952	1.6e-02s, 10	1e-03 s, 5.2 %

The table below Fig. 4 shows the means and variances for the Gaussian prior, the true posterior and the approximate Gaussian posteriors following the discrete fusion updates. 10 trials of LW with 500 samples were used to generate the statistics (mean and std. dev.) for the LW posterior estimates, and the best LW solution over all trials is plotted in blue for each case in Fig. 4. The table also gives the times (in Matlab) needed to run each approximate fusion update and relevant measures of computational efficiency. For VB, the number of EM steps for a converegence tolerance of 1e-3 on (18) is shown (starting from the same random initialization of α and ξ_c in each case); for LW, the sampling efficiency (SE) is shown (defined by the effective sample size divided by the number of samples) [15]. The SE is a figure of merit for a Monte Carlo estimator: larger SE means that samples are representative of the true posterior, so that subsequent sample-based estimates are more reliable.

From the table, we see that the VB posterior mean is very close to the true posterior mean in all cases, and that the VB computations are real-time (under 0.02 secs in all cases). VB's robustness is demonstrated by the fact that, even though the observations in cases (a)-(c) are progressively more 'surprising' with respect to the prior, the closeness of the VB mean to the true posterior mean is always maintained. We also see that the number of EM steps needed in each case is small, so that convergence to the unique posterior approximation is quick. The LW Gaussian posterior does well in case(a) since the prior 'agrees' with the discrete observation. However, the SE in this case is only 50%, which means only half of the 500 samples drawn in the 1-D space are contributing to estimate. LW starts to show much lower SE (and larger estimation variance) in the cases (b) and (c). The results are particularly bad in (c), as the best LW Gaussian posterior is off in the tail of the true posterior distribution. In contrast, VB shifts towards the high probability mass region of the true posterior. Notice that in all three cases, the variational estimate of the posterior variance is consistently optimistic; this is typical of VB methods [9].

C. 2D Localization Scenario with Continuous Measurements

Now consider a scenario where the robot moves deterministically to the right with constant velocity starting from $Z = [1,1]^T$, moving 1.5 m per time step. The bearing sensor now takes measurements to the target with zero mean Gaussian noise with variance 0.25 rad^2, and the target's true location is $x = 5$ m, $y = 5.7$ m. The robot takes 3 steps with 1 measurement each and fuses them with a diffuse Gaussian target prior to produce posterior $p(X|B_{1:3}, Z_{1:3})$. Non-linear least squares (NLS) is used for nearly optimal continuous fusion to give a Gaussian posterior with mean μ_{NLS} at the NLS solution and covariance matrix Σ_{NLS} based on the NLS Hessian at μ_{NLS} [6]. The human provides no observations for k=1 or 2; at k=3, the human provides one of $m = 24$ possible observations $D(k) \in \mathcal{D} = \{$'0-2 m', '2-4 m', '$>4$ m'$\} \times \{$'North', 'NorthWest', 'West', ..., 'NorthEast'$\}$, whose likelihood is modeled by (2) with $Q = [Z(k), X]^T$. The realization used is shown in Figure 5; the whitespace corresponds to Q where $p(D = j|Q) \approx 1$ for some j (e.g. $(x_{rel}, y_{rel}) = (0, -3)$ is where $D(k)$ is most likely '2-4 m South'). To account for the fact that $D(k)$ gives *relative* range and bearing information, the softmax biases for each class c are adjusted via $b'_c = b_c - w_c^T Z(k)$ so that $p(D|Q)$

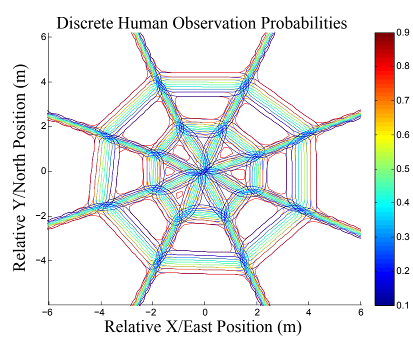

Fig. 5. Softmax contours of $p(D = j|Z(k), X)$ for $m = 24$ discrete human observations j of range and bearing as a function of $X - Z(k)$.

always centers on $Z(k)$.

We examine the fusion of $D(3)$ =‘$>$ 4 m North’ for time step k=3 with ‘prior’ $p(X|B_{1:3}, Z_{1:3})$. The black cross and 2σ ellipse in Fig. 6 shows μ_{NLS} and Σ_{NLS} for $p(X|B_{1:3}, Z_{1:3})$, and the colored contours show the likelihood for $D(3)$ =‘$>$ 4 m North’ from Fig. 5 (the true X and $Z(k)$ are also shown). The ‘best Gaussian’ approximation to the true fully fused posterior $p(X|D(3), B_{1:3}, Z_{1:3})$ is shown in Fig. 6 by the blue ellipsoid, where ‘exact’ posterior mean μ_* and covariance Σ_* are shown. These parameters were estimated using 5000 runs of random-walk Markov Chain Monte Carlo sampling, which required 5 secs to run. This represents the best Gaussian approximation to $p(X|D(3), B_{1:3}, Z_{1:3})$ in the ‘moment-matching’ sense [9]. We use this here to show how closely the VB and LW fusions come to matching μ_* and Σ_*. It is clear from the best Gaussian approximation that fusion of $D(3)$ improves the NLS estimate: μ_* is closer to the true X and $|\Sigma_*| \leq |\Sigma_{NLS}|$.

The VB and LW fusions are shown in Fig. 6 by the magenta and green Gaussian ellipsoids, respectively. 10 LW trials were run with 100 samples to produce average statistics for the μ and Σ estimates, and the best LW Gaussian estimate (with smallest $||\mu - \mu_*||^2$) is shown in the plot. VB and LW both have μ very close to μ_*, and both methods are optimistic with respect to Σ_*. VB achieved convergence to the same previous tolerance with 18 EM steps in 0.12 secs; this longer time arises since m is larger. LW took an average of 1e-03 secs but only had SE of about 15%, since μ_{NLS} and Σ_{NLS} ‘disagree’ with $D(3)$.

V. Conclusions and Future Work

We have derived and demonstrated the properties and utility of a new technique for VB fusion of multi-class discrete observations of continuous states with conventional continuous sensor measurements. This method produces fast closed-form Gaussian posterior approximations and has several nice properties that make it suitable for online hybrid inference. The VB method has immediate real-time applications to cooperative HR estimation and is easily applicable to other HRI problems. We have also been able to extend the theory of Section III to handle Gaussian mixture priors and softmax mixture likelihood functions, so that approximate VB posteriors are guaranteed to be unique Gaussian mixtures in more complex settings. These extensions also enable robust VB approximate inference in general hybrid Bayesian networks via the junction tree method, as described in [8] for the binary VB method; these extensions are not given here due to limited space.

We are working on fixing the VB fusion’s optimistic posterior covariance, which can lead to inconsistent estimates. As [16] notes for the binary VB method, the closeness of the VB mean to the true posterior mean can be used to refine the covariance (e.g. via post-hoc sampling). As Section IV mentions, the VB method could also be used in the HR localization problem when the robot location is uncertain. For this, we intend to use the VB method to obtain Rao-Blackwellized importance sampling estimates [15] of target location based on fusion of discrete human observations, robot bearing sensor data, and robot pose information.

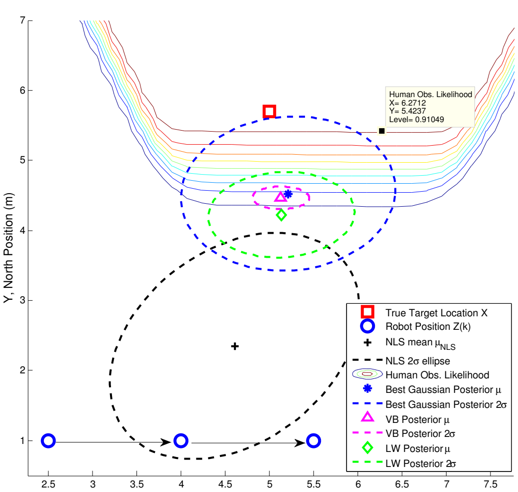

Fig. 6. Setup and results for 2D hybrid fusion scenario.

Acknowledgments: This work is supported by a National Science Foundation Graduate Research Fellowship and AFOSR contract number FA9550-08-0356.

References

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

[2] R. Jones, E. Connors, and M. Endsley, “Incorporating the human analyst into the data fusion process by modeling situation awareness using fuzzy cognition maps,” in *FUSION 2009*, pp. 1265–1271.

[3] T. Kaupp, A. Makaerenko, F. Ramos, B. Upcroft, S. Williams, and H. Durrant-Whyte, “Adaptive human sensor model in sensor networks,” in *FUSION 2005*, vol. 1, pp. 748–755.

[4] T. Kaupp, A. Makarenko, S. Kumar, B. Upcroft, and S. Williams, “Operators as information sources in sensor networks,” in *IROS 2005*, pp. 936–941.

[5] 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.

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

[7] U. Lerner, “Hybrid bayesian networks for reasoning about complex systems,” Ph.D. dissertation, Stanford University, 2002.

[8] K. Murphy, “A variational approximation for bayesian networks with discrete and continuous latent variables,” in *Proc. UAI 1999*.

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

[10] G. Bouchard, “Efficient bounds for the softmax function and applications to approximate inference in hybrid models,” in *NIPS*, vol. 21.

[11] S. Osentoski, V. Manfredi, and S. Mahadevan, “Learning hierarchical models of activity,” in *IROS 2004*, pp. 891–896.

[12] O. Schrempf, D. Albrecht, and U. Hanebeck, “Tractable probabilistic models for intention recognition based on expert knowledge,” in *IROS 2007*, San Diego, CA, pp. 1429–1434.

[13] P. Krauthausen and U. Hanebeck, “Intention recognition for partial-order plans using dynamic bayesian networks,” in *FUSION 2009*, pp. 444–451.

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

[15] J. Liu, *Monte Carlo Strategies in Scientific Computing*. New York: Springer, 2001.

[16] T. Jaakkola and M. Jordan, “Bayesian parameter estimation via variational methods,” *Statsistics and Computing*, vol. 10, pp. 25–37, 2000.

[17] J. Schoenberg, M. Campbell, and I. Miller, “Localization with multi-modal vision measurements in limited GPS environments using Gauss-sum filters,” in *ICRA 2009*, pp. 1423–1428.