Proceedings of IMECE 2007
ASME 2007 International Mechanical Engineering Congress and Exposition
November 11-15, 2007, Seattle, USA

IMECE2007-43791

INFORMATION-THEORETIC SENSOR MOTION CONTROL FOR DISTRIBUTED ESTIMATION

Allison D. Ryan*
Vehicle Dynamics Lab
University of California
Berkeley, California 94720
allison@vehicle.me.berkeley.edu

Hugh Durrant-Whyte
ARC Centre of Excellence
for Autonomous Systems
The University of Sydney
Sydney, NSW
Australia

J. Karl Hedrick
Vehicle Dynamics Lab
Department of Mechanical Engineering
University of California
Berkeley, California, 94720

ABSTRACT

Estimate uncertainty is a clear metric for sensing problems, but is not traditionally used for optimal control of mobile sensors because it is difficult to model how it is affected by sensor motion. This work develops a multiple-step receding horizon cost for sensor motion control based on minimization of expected entropy of the estimate distribution. The structure of the cost function is analyzed and used to upper bound the degree of coupling between sensors. The contribution is a multiple step prediction of the estimate entropy incorporating probabilistic sensor and target motion models and a decomposition for its decentralized calculation. Multiple-step receding horizon control has the potential to provide better performance than single-step optimization in the cases of delayed payoff or sensor motion constraints, and has not generally been implemented for non-Gaussian models. An example is developed based on a team of unmanned air vehicles carrying vision sensors, and initial simulation results confirm the accuracy of the predictive cost calculation.

NOMENCLATURE

$P(x)$	Distribution of the state estimate
$H(P(x\|z))$	Entropy of conditional state estimate
$\mathbf{H}(P(x\|z))$	$\int_z P(z)H(P(x\|z))$, the expected conditional entropy
$I(x;z)$	The mutual information of random variables x and z
u_k^j	control of agent j at time k
z	observations, same subscript/superscript notation as u
$z^{\sim q}$	observations from all agents other than q
R^i	The region of x that is observed by sensor i
$R^i - R^j$	The region of x that is observed by sensor i but not sensor j
$x \sim \mathcal{N}(x;\mu,\sigma)$	x is a Gaussian random variable with mean μ and variance σ

INTRODUCTION

Distributed mobile sensing platforms such as uninhabited air vehicles (UAVs), underwater vehicles (UUVs), and ground vehicles (UGVs) can collect information over wide and inaccessible areas. Computing, sensing and communication technology have developed to the point where autonomous mobile sensors can be built and networked using mostly off the shelf components, but the development of robust and efficient collaborative motion control algorithms for collecting desired information is an ongoing area of research.

The contribution of this paper is a receding horizon control formulation for a team of sensors to cooperatively search for and track a target. This formulation incorporates probabilistic sensor and target motion models and sensor motion constraints and explicitly optimizes the quality of the estimate. An upper bound on the coupling between sensors is derived which can be efficiently calculated by one sensor without a full model of the other.

The simplest interpretation of the mobile sensing problem is to treat it as a sensor delivery task. The result is a multi-vehicle

*Address all correspondence to this author.

routing problem [1], which can be extended to include ideas such as heterogeneous sensors and prioritized tasks. A flight-tested example of this type of UAV sensing system has been implemented by the authors and will provide a platform for future testing [2].

Greater autonomy is provided when a mission is described in terms of the desired information to be collected, without requiring that a location be known *a priori*. Applications such as search, mapping and target tracking using teams of mobile sensors require motion control to optimize the acquisition of useful observations. Control and estimation cannot be considered independently as in classical methods such as linear quadratic Gaussian control because the availability of observations depends on the sensor motion.

Traditionally, search techniques attempt to maximize the probability of detecting a target [3] and target tracking is often accomplished by commanding a feasible sensor trajectory defined relative to the estimated target position [4]. Searching and tracking are rarely addressed in the same framework in spite of the fact that the transition between the two is often unclear. For targets with uncertain motion and sensors with constrained motion and narrow field of view, the target may be observed intermittently, suggesting the need for a single control formulation based explicitly on the quality of the target state estimate.

Existing methods for combined search and track are often ad-hoc. Webb and Furukawa define a search and track problem using probability of detection as a value function and a consistent sensor model that includes the possibility of not detecting the target [5]. However, probability of detection does not capture the idea that some observations are more informative than others. For example, a bearing-only sensor must observe the target from more than one angle in order to fully localize it.

Much of the previous work on information-theoretic control has been based on "information surfing" or receding horizon optimization of information gain over a single step [6–9]. However, receding horizon optimization with multiple-step lookahead has the potential to produce far better solutions by accounting for sensor motion constraints and delayed payoff, although it comes at a high cost of computation. In previous solutions with a longer optimization horizon, either the target state is assumed constant or closed form estimation such as a Kalman filter provides multiple-step prediction [10].

The paper will be organized as follows. First the sensor motion control problem is formulated as a minimization of the expected estimate entropy. Selection of entropy as a cost function gives the problem a useful structure which is suited to decentralization. Methods for predicting the entropy resulting from a sequence of future observations are derived. An example is developed based on UAVs mounted with bearing-only sensors, and finally contributions and future work are discussed.

THE ESTIMATION PROBLEM

The goal in any estimation problem is to reduce the uncertainty of the the estimate of a target state x. Conditional entropy is a measure of the uncertainty when the estimate is conditioned on a set of observations. To predict the conditional entropy that will result from future observations, we take the expectation over the values of those observations. Thus, the expected uncertainty of the estimate after making observations z is represented by the expected conditional entropy $\mathbf{H}(P(x|z))$. It should be noted that in this formulation, x and z are both random variables whose values are unknown, although z will eventually be observed.

$$\begin{aligned} H(P(x)) &= -\int_x P(x)\log P(x)dx \qquad (1)\\ \mathbf{H}(P(x|z)) &= \int_z P(z)H(P(x|z))dz \end{aligned}$$

It will be useful to write the entropy of x conditioned on z in terms of mutual information, $I(x;z)$. Equation (2) defines the mutual information of two random variables in terms of the expected reduction in uncertainty of one variable after observing the other.

$$\begin{aligned} I(x;z) &= H(P(x)) - \mathbf{H}(P(x|z)) \qquad (2)\\ &= \int_x\int_z P(x,z)\log\frac{P(x,z)}{P(x)P(z)}dz\,dx \end{aligned}$$

Although the goal of the estimation task is to minimize uncertainty over the duration of the task, a cost function cannot generally be evaluated over such an indeterminant duration except under limiting assumptions. Therefore the infinite horizon minimization is reduced to a finite horizon problem in Eqn. (3) by defining the cost J as the sum of the expected conditional entropy from the current time k up to horizon time $k+N$. At each time k, an optimal control sequence $u_{k:k+N}$ is computed and only the first control u_k is applied. This receding horizon optimization does not guarantee optimality over the infinite horizon, but is a generally accepted method [11].

$$J = \sum_{i=k}^{k+N}\mathbf{H}(P(x_i|z_k,\ldots,z_i)) \qquad (3)$$

A team of M sensors generates a set of simultaneous observations $z_k^1,\ldots,z_k^M$ at each time k. Equation (4) shows how the conditional entropy of x given the set of simultaneous observations breaks down into mutual information and unconditional entropy terms.

$$\mathbf{H}(P(x|z^1,\ldots,z^M)) =$$
$$H(P(x)) - \sum_{j=1}^{M} \left(I(x;z^j) + H(P(z^j))\right) - H(P(z^1,\ldots,z^M)) \quad (4)$$

Equation (5) shows how the expected conditional entropy (4) decomposes into a constant C, individual agent terms J_j and a coupled term J_c. The coupled term includes the joint entropy of the team's full set of measurements.

$$J = C + J_c(u^1,\ldots,u^M) + \sum_{j=1}^{M} J_j(u^j) \quad (5)$$
$$C = (1-M)\sum_{i=k}^{k+N} H(P(x_i))$$
$$J_j(u^j) = \sum_{i=k}^{k+N} \mathbf{H}(P(x_i|z^j_{1:i}))$$
$$J_c(u^1,\ldots,u^M) = \sum_{i=k}^{k+N} \left(\sum_{j=1}^{M} H(P(z^j_{1:i}))\right) - H(P(z^1_{1:i},\ldots,z^M_{1:i}))$$

The coupling term represents the redundancy of the information gain between sensors. A control selected to optimize only the individual costs (neglecting J_c) could result in wasted effort from all sensors but one if all measurements after the first provide the same information.

STRUCTURE OF INFORMATION-THEORETIC COST FUNCTION

The team goal is to choose the set of individual controls $\{u^j\}$ to minimize the cost J. A stationary point of J is characterized by the gradient condition $\nabla J = 0$ where the gradient is taken with respect to the set of control variables. If the stationary point also satisfies the second derivative condition, $\nabla^2 J \succ 0$, then it is a local minimum. Gradient-based minimization methods incrementally improve a solution by moving it away from the gradient direction and terminate when a stationary point is achieved. The body of literature on these methods is extensive [12].

Gradient-based optimal control in distributed systems requires that various components of the control be calculated on separate processors, often under communication constraints. With structure of the cost function given by Eqn. (5), the terms $J_j(u_j)$ each depend on only the individual control component, but the coupling term $J_c(u^1,\ldots,u^M)$ depends on the control decisions of all agents. Therefore the gradient condition breaks down into gradients with respect to each agent's control as follows.

$$\nabla J = 0 \Leftrightarrow \nabla_j J = 0 \;\forall j \in [1,M] \quad (6)$$
$$\nabla_j J = \frac{\partial}{\partial u^j}\left(J_j(u^j) + J_c(u^1,\ldots,u^M)\right)$$

Decentralized implementation of gradient-based optimization is an active area of research that is outside the scope of this paper [13]. In general, information from each processor is required to calculate a single gradient step in the optimization, leading to high communication requirements. Mathews et. al. discuss the effect of communication frequency, delays, and degree of coupling [14] in this type of problem. Although decentralized gradient descent will converge to a team local minimum (under conditions such as convexity and continuity), it requires a full round of communication in order to calculate each gradient step, and many steps may be required to achieve the a stationary point with the desired tolerance. This communication rate may not be achievable for mobile sensing applications due to limited bandwidth and connectivity.

The coupling between individual agents' effects on the cost governs the necessary degree of cooperation. For example, if two sensors i and j measure features of the environment which are fully independent, the observations z^i and z^j are uncorrelated, and so are the agents' contributions to the cost, resulting in Eqn. (7). In this case, each agent can calculate the gradient of the cost with respect to its own control without knowledge of the control of the other agent.

$$\frac{\partial}{\partial u^i \partial u^j} J_c = \frac{\partial}{\partial u^j \partial u^i} J_c = 0 \quad (7)$$

An agent can choose a decentralized optimization strategy for cooperation by estimating the degree of its coupling with other agents. The coupling term $J_c(u^1,\ldots,u^M)$ depends on the joint probability of all of the agents' observations, $P(z^1,\ldots,z^M)$. Calculation of the full joint distribution is generally intractable, and does not provide direct insight into the degree of coupling between sets of agents.

For a single agent q to evaluate the team coupling, it is useful to decompose $J_c(u^1,\ldots,u^M)$ from Eqn. (5) into terms due to z^q and terms due to the observations from all other agents, $z^{\sim q} \equiv \{z^1,\ldots,z^{q-1},z^{q+1},\ldots,z^M\}$. This is shown in Eqn. (8).

$$J_c = \sum_{i=k}^{k+N} \sum_{m \neq q} H(P(z^m_{1:i})) - H(P(z^{\sim q}_{1:i})) + I(z^q_{1:i}; z^{\sim q}_{1:i}) \quad (8)$$

Agent q can further decompose the team coupling by approximating $I(z^j; z^{\sim j})$ by pairs of agents as in Eqn. (9). This corresponds to assuming that no more than two sensors observe the "same" information, and may over or under estimate J_c in the presence of three-way (or greater) coupling.

$$I(z^j; z^{\sim j}) \approx \sum_{\substack{m=1 \\ m \neq j}}^{M} I(z^j; z^m) \tag{9}$$

Equation (9) allows agent q to approximate the total team coupling in terms of its coupling with each other agent individually.

$$J_c \approx \sum_{i=k}^{k+N} \sum_{\substack{m=1 \\ m \neq q}}^{M} H(P(z^m_{1:i})) - H(P(z^{\sim q}_{1:i})) + \sum_{\substack{m=1 \\ m \neq q}}^{M} I(z^q_{1:i}; z^m_{1:i}) \tag{10}$$

CALCULATION OF PREDICTIVE INFORMATION-THEORETIC COST FUNCTION

Receding horizon optimization based on an information-theoretic cost function as proposed in the previous sections requires the capability to estimate the entropy of the estimate distribution conditioned on observations which have not yet occurred. For single step lookahead, the conditional entropy after one observation can be estimated in two steps. First, the current estimate is convolved with the motion model to produce a predicted estimate. Then, the mutual information of the future observation and the predicted estimate is calculated from Eqn. 2. The expected conditional entropy is equal to the entropy of the predicted estimate minus the mutual information. Unfortunately, this method does not extend to greater than one step lookahead because the effect of a general probabilistic motion model can only be calculated for a particular distribution and there is not a convenient method to average over the future observations.

Receding horizon control with multiple step lookahead requires estimation of the information gain due to a series of observations. Even for a single sensor, this calculation becomes difficult in the presence of target motion except in the linear Gaussian case. Local estimation of the effect of a series of observations from a set of distributed sensors requires a variety of approximate methods depending on the degree of coupling between observations.

In spite of the high cost of computation, multiple step receding horizon control has the potential to provide high payoff for UAV applications by allowing the UAV to plan motion on a scale at least sufficient to model its turn radius. In scenarios with obstacles, the motion planning horizon *must* be sufficient for the UAV to execute an avoidance maneuver [15]. Figure 1 compares two different planning horizons for a bearing-only sensor observing a slowly moving target. The trajectory calculated using a receding horizon control with horizon length $N = 4$ observes the

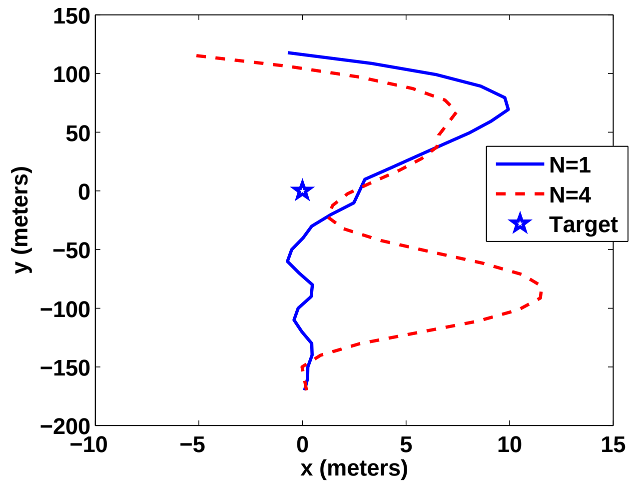

Figure 1. UAV PATH FOR BEARING-ONLY SENSOR AND SLOWLY VARYING RANDOM WALK TARGET. ENTROPY-MINIMIZING CONTROL RESULTS IN WIDER SWEEP PATTERN FOR LONGER OPTIMIZATION HORIZON.

target from widely varying orientations, whereas the single-step optimal planner produces a more direct trajectory.

It is well known that differently oriented observations from a bearing-only sensor are more useful than parallel observations, so the control with longer horizon will be more beneficial for localization. Figure 2 illustrates this idea by showing the combined effect of two likelihood functions shown in contour: one case (A) where the measurements are nearly parallel, and another case (B) where they are not. The dark shaded area represents the combination of the two likelihood functions, and the case B reduces to a smaller area. Figures 1 and 2 reflect the same scenario described in the tracking application and simulations section.

Individual information gain

Recall that the team cost in Eqn. (3) is equal to the expected estimate entropy, summed at each time from the present up to the receding horizon. The cost is broken down into individual components which can be calculated with information from a single sensor, and a coupling term which depends on all sensors. The individual cost component for an agent j is repeated from Eqn. (5).

$$J_j = \sum_{i=k}^{k+N} \mathbf{H}(P(x_i | z^j_{1:i})) \tag{11}$$

For estimation of a constant state, such as localization of a stationary feature, representation of expected entropy simplifies

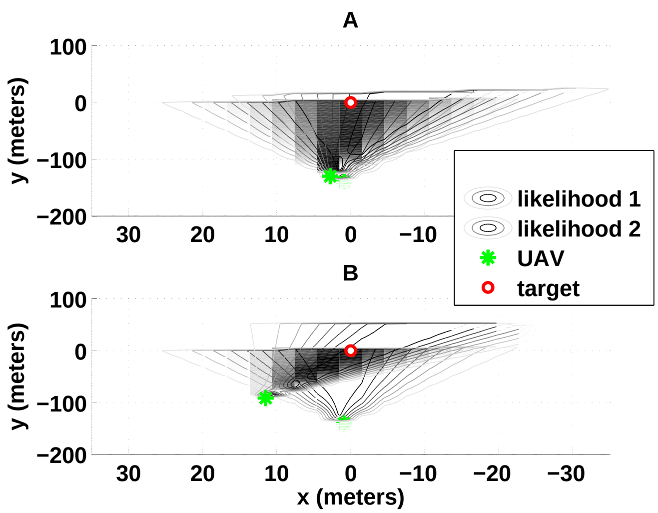

Figure 2. COMBINATION OF LIKELIHOOD FUNCTIONS FROM BEARING ONLY SENSORS. CONTOUR LINES SHOW SUCCESSIVE LIKELIHOOD FUNCTIONS AND SHADING REPRESENTS COMBINED LIKELIHOOD (PRODUCT).

because the distribution of x changes only due to observations,

$$P(x_N|z_1,\ldots,z_{i-1}) = P(x_i|z_1,\ldots,z_{i-1}) \tag{12}$$

leading to the following result.

$$\mathbf{H}(P(x_N|z_1,\ldots,z_N)) = H(P(x_N)) - \sum_{i=1}^{N} I(x_N;z_i) + \sum_{i=2}^{N} I(z_i;[z_1\ldots z_{i-1}]) \tag{13}$$

Thus, multi-step optimization for estimation of a constant state can be performed based only on calculations of information gain at the time of individual measurements ($I(x_i;z_i)$) and the redundancy of the measurements ($I(z_i;z_j)$).

When the target state evolves, the cost calculation is an expectation over the sequence of future observations and cannot be written only in terms of mutual information because the target motion model must be included.

$$\mathbf{H}(P(x_N|z_1,\ldots,z_N)) = \int_{z_1}\cdots\int_{z_N} P(z_1,\ldots,z_N)H(P(x_N|z_1,\ldots,z_N))dz_1\ldots dz_N \tag{14}$$

In most cases, the effect of target motion on the estimate $P(x)$ can only be determined by convolving the motion model with the previous estimate. This means that in the presence of non-linear target motion or non-Gaussian distributions, a closed form expression for the conditional entropy such as Eqn. (13) cannot be obtained.

The joint distribution $P(z_1,\ldots,z_N)$ depends on the effect of the process model at each step as in Eqn. (15), and so can only be approximated by simulation for the general non-linear non-Gaussian case. For a scalar sensor discretized to m possible outputs, the size of the joint distribution is m^N. Even for a binary sensor, calculation of the joint distribution is not likely to be feasible in real time due to the complexity of calculating the probability of a single sequence of observations. Therefore the expectation in Eqn. (14) is approximated by sampling from $P(z_1,\ldots,z_N)$.

$$P(z_1,\ldots,z_N) = P(z_1)\prod_{i=2}^{N} P(z_i|z_1,\ldots,z_{i-1}) \tag{15}$$
$$= \int_{x_1} P(z_1|x_1)P(x_1)dx_1 \prod_{i=2}^{N}\int_{x_i} P(z_i|x_i)P(x_i|z_1\ldots z_{i-1})dx_i$$

With a method to calculate $J_j(u^j)$, an individually optimal solution $u^{j*} = \arg\min J_j(u^j)$ can be calculated using a gradient-based method. Unless $J_j(u^j)$ is a convex function of u^j, the gradient-based method will converge to a local optimum, and so will be sensitive to the initial condition of the algorithm. Figure 3 shows two trajectories for a bearing-only sensor which are optimal solutions to Eqn. (3), and one suboptimal solution that exists "between" them in the control space. This example shows that the proposed cost function is not convex.

To encourage the optimization to converge to a desirable local minimum, an initial condition can be generated by a proportional heading rate control which steers the sensor toward the most likely target location. This is motivated by the intuition that observing the target will be more informative than not observing it, and sensors are generally more informative at close range.

Estimation of pairwise coupling

In addition to the information gain due to each individual sensor, the information-theoretic cost depends on the coupling between the set of observations. Recall that any agent q can approximate its coupling with the rest of the team using Eqn. (10), repeated below.

$$J_c \approx \sum_{i=k}^{k+N}\sum_{\substack{m=1\\m\neq q}}^{M} H(P(z^m_{1:i})) - H(P(z^{\sim q}_{1:i})) + \sum_{\substack{m=1\\m\neq q}}^{M} I(z^q_{1:i};z^m_{1:i}) \tag{16}$$

For agent q to select its individual control u^q which maximizes the total cost, it need only calculate terms which depend on

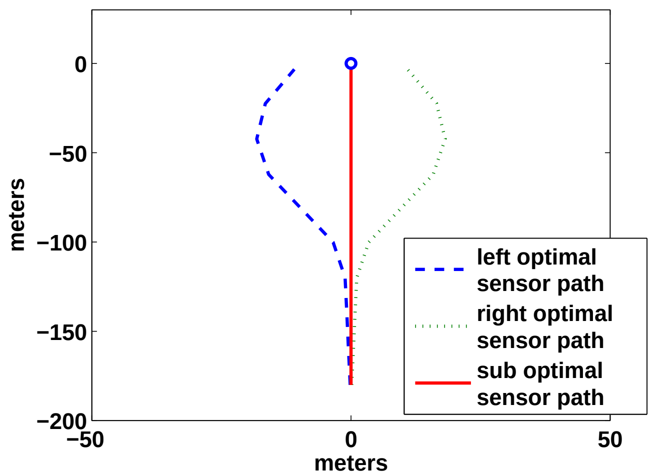

Figure 3. PATHS OF BEARING-ONLY SENSORS WITH LIMITED TURN RATE OBSERVING STATIONARY TARGET. LOCALLY OPTIMAL SOLUTIONS ARE SYMMETRIC AND SEPARATED BY A LOCALLY MAXIMUM COST PATH.

u^q. In gradient-based optimization techniques, an optimal control results from a series of iterations that depend on the gradient of the cost with respect to the control, $\frac{\partial}{\partial u^q}J$. Equation (17) shows the elimination of non-local terms. The coupling is now reduced to pairwise terms $I(z^q_{k:k+N}; z^m_{k:k+N})$, based on the pairwise coupling approximation discussed previously.

$$\frac{\partial}{\partial u^q}J = \frac{\partial}{\partial u^q}J_q(u^q) + \sum_{\substack{m=1\\m\neq q}}^{M} \frac{\partial}{\partial u^q} I(z^q_{1:i}; z^m_{1:i}) \tag{17}$$

The mutual information between two simultaneous observations depends on the underlying state estimate, requiring an integral over the state space as well as the observations as shown in Eqn. (18). An upper bound for $I(z^i, z^j)$ that does not require integration over the state space is desirable because the state space is generally much larger than the pairwise observation space, and the integral can not generally be evaluated analytically.

$$\begin{aligned} I(z^i, z^j) &= \int_{z^i}\int_{z^j} P(z^i, z^j) \log \frac{P(z^i, z^j)}{P(z^i)P(z^j)} dz^i\, dz^j \\ P(z^i, z^j) &= \int_x P(z^i|x)P(z^j|x)P(x)dx \end{aligned} \tag{18}$$

Due to the difficulty of calculating mutual information directly from Eqn. (18), an upper bound will be derived that does not require the joint distribution $P(z^i; z^j)$. It is useful to note that most sensors have a limited range, and the degree of coupling between the observations from two sensors is strongly related to the degree to which their observed areas, or footprints, overlap. If the target is not in the footprint R^j of sensor j, the observation will be $z^j = \emptyset$, or “no observation” (assuming no false detections).

The observation z^j from a scalar sensor is thus in the set $\mathbb{R} \bigcup \emptyset^j$ where $\emptyset^j$ is defined as the lack of an observation from sensor j. This type of sensor model which includes the possibility of not detecting the target provides a single estimation model (and therefore cost function) for searching and tracking applications. The probability density function of z must still integrate to one, but the integral includes the value $z = \emptyset$. Therefore the integral over over values of $z \in \mathbb{R}$ must equal the probability of detection, or $1 - P(z = \emptyset)$.

For example, consider a range measurement with Gaussian noise, given by Eqn. (19). Range is measured from the sensor position $\mathbf{s}$ to the target position $\mathbf{x}$, with additive Gaussian noise with variance σ. This model can easily be extended to include a probability of missed or false detection as well. The range and probabilities of missed or false detections form the *detection model* of the sensor.

	$x \in R^j$	$x \notin R^j$
$z \in \mathbb{R}$	$P(z\|x) = \mathcal{N}(z; \|\|s - \mathbf{x}\|\|, \sigma)$	$P(z \in \mathbb{R}\|x) = 0$
$z = \emptyset$	$P(z = \emptyset\|x) = 0$	$P(z = \emptyset\|x) = 1$

(19)

In the case where both sensors have limited ranges which *do not* overlap ($R^i \cap R^j = \emptyset$), $I(z^i; z^j)$ simplifies to Eqn. (20) and depends only on the target state estimate and the detection models of each sensor.

$$\begin{aligned} I(z^i; z^j) &= (P(\emptyset^i) - 1)\log P(\emptyset^j) + (P(\emptyset^j) - 1)\log P(\emptyset^i) \\ &+ P(\emptyset^i, \emptyset^j) \log \frac{P(\emptyset^i, \emptyset^j)}{P(\emptyset^i)P(\emptyset^j)} \end{aligned} \tag{20}$$

An upper bound for $I(z^1; z^2)$ is derived in Eqn. (21) from the definition by lower bounding the entropy of z^2 conditioned on z^1. The lower bound on conditional entropy can be derived intuitively from a “most informative, limited range” model of sensor one. The most that sensor one can reduce the entropy of $P(z^2)$ is if z^1 provides ”perfect information” within the limits of its range. Perfect information implies that z^1 is an accurate measurement of x with no uncertainty, if $x \in R^1$.

$$\begin{aligned} I(z^1; z^2) &= H(P(z^2)) - \mathbf{H}(P(z^2|z^1)) \\ &< H(P(z^2)) - \min \mathbf{H}(P(z^2|z^1)) \end{aligned} \tag{21}$$

Under the "most informative, limited range" model of sensor one given in Eqn. (22), $z^1 = x$ when the target position x is in range of sensor one. If $x \in R^1 - R^2$ then the target is outside the footprint of sensor two but inside the footprint of sensor one, and $z^2 = \emptyset^2$ (no uncertainty). If $x \in R^1 \cap R^2$ then $z^1 = x$ and z^2 will be distributed about x according to the model of sensor two. If $z^1 = \emptyset^1$ the target is known to be outside of R^1 and $P(z^2|z^1)$ is equal to $P(z^2|x \notin R^1)$.

$$\begin{array}{c|c} z^1 & P(z^2|z^1) \\ \hline z^1 = x \in R^1 - R^2 & z^2 = \emptyset^2 \\ z^1 = x \in R^1 \cap R^2 & P(z^2|z^1) = P(z^2|x) \text{ (sensor model)} \\ z^1 = \emptyset^1 & P(z^2|z^1) = \frac{1}{P(\emptyset^1)} \int_{x \notin R^1} P(z^2|x)P(x)dx \end{array} \quad (22)$$

An expected conditional entropy for the most informative sensor is calculated based on these three possible outcomes and given by Eqn. (23). The mutual information of z^2 with the most informative sensor model for z^1 will be referred to as the "most informative upper bound" of $I(z^1;z^2)$, or $\hat{I}(z^1;z^2)$. Sensor two can calculate $\hat{I}(z^1;z^2)$ knowing only R^1 rather than the full model of sensor one. Similarly, sensor one can calculate a most informative upper bound based on R^2. Even in heterogeneous teams, sensors can upper bound their coupling by communicating sensing regions rather than full sensor models. In general, the value of $\hat{I}(z^1;z^2)$ will depend on which sensor is approximated using the most informative model and which sensor is modeled more completely.

$$\begin{aligned} &\min \mathbf{H}(P(z^2|z^1)) = \\ &\quad P(x \in R^1 \cap R^2)\ H(P(z^2|x)) + P(x \notin R^1)\ H(P(z^2|\emptyset^1)) \end{aligned} \quad (23)$$

Figures 4 and 5 compare the most informative upper bound to $I(z^1;z^2)$ for two different distributions $P(x)$. In both cases, x, z^1 and z^2 are scalar, with Gaussian sensor noise ($P(z^i|x) \sim \mathcal{N}(z^i;x,\sigma)$). The target state x is uniformly distributed in Fig. 4 and Gaussian distributed in Fig. 5.

TRACKING APPLICATION AND SIMULATIONS

The algorithms developed in the preceding sections are applied to an example of two fixed-wing UAVs carrying bearing-only sensors. Monocular vision can be considered a bearing-only sensor in the absence of additional information (such as target size) used to estimate range. A single target is detected and tracked, but the UAV motion constraints prevent it from being continuously observed.

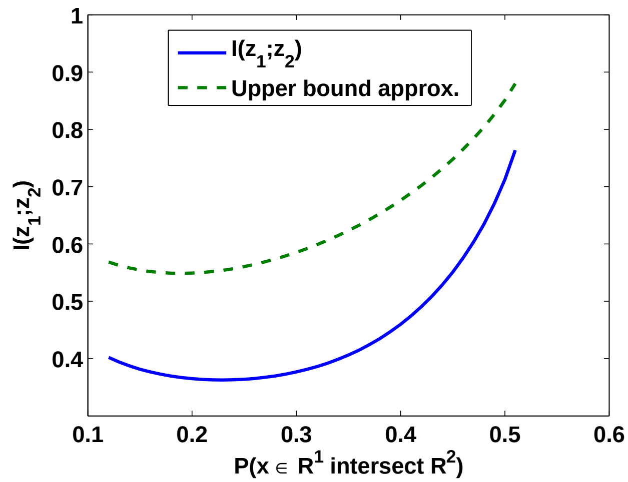

Figure 4. MUTUAL INFORMATION AND UPPER BOUND APPROXIMATION FOR TWO SENSORS WITH OVERLAPPING RANGES, CONDITIONED ON UNIFORM $P(x)$.

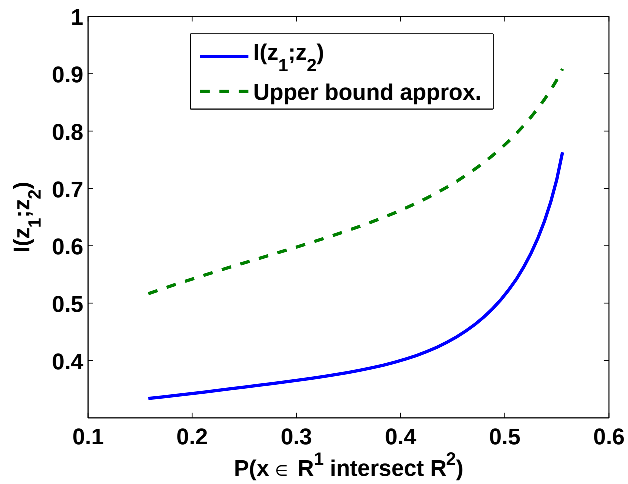

Figure 5. MUTUAL INFORMATION AND UPPER BOUND APPROXIMATION FOR TWO SENSORS WITH OVERLAPPING RANGES, CONDITIONED ON GAUSSIAN $P(x)$.

Process and sensor models

The fixed wing UAV motion is represented by a constant altitude, constant airspeed (V), bounded turn rate kinematic model: Eqn. (24). A roll angle ϕ is modeled from the turn rate ψ based on the coordinated turn model in Eqn. (25). This model is commonly used for high-level control with the assumption of an autopilot to provide stability and the desired turn rate [16]. The target motion is modeled as a Gaussian random walk with variance 10 meters, and the simulation time step is one second.

UAV	velocity	20 m/s
	$\dot{\psi}_{max}$	0.2 r/s
	altitude	100 m
Sensor	view angle	$\pi/4$ r
	P_{miss}	0.1
	σ	0.05 r

Table 1. Simulation parameters

$$\begin{aligned}\dot{x} &= V\cos\psi \\ \dot{y} &= V\sin\psi \\ \dot{\psi} &= u \\ |u| &\le \dot{\psi}_{max}\end{aligned} \tag{24}$$

$$\begin{aligned}\tan\phi &= \frac{\text{radial acceleration}}{\text{gravity}} \\ &= \frac{V\dot{\psi}}{g}\end{aligned} \tag{25}$$

The bearing-only sensor model corresponds to a fixed camera mounted below the UAV. Its limited view angle and the UAV's altitude and roll angle define a sensor footprint at ground level, which is represented by the sensor region R^j in previous sections. The observation is modeled as Eqn. (26). It is Gaussian distributed about θ if the target is within R^j and will not be observed ($z = \emptyset$) otherwise. Simulation parameters for the UAV and sensor are given in Table 1.

	$x \in R^j$	$x \notin R^j$
$z \in [0, 2\pi]$	$P(z\|x) = (1 - P_{miss})\mathcal{N}(z;\theta,\sigma)$	$P(z\|x = 0$
$z = \emptyset$	$P(z = \|x) = P_{miss}$	$P(z\|x) = 1$

(26)

Recursive Bayesian filtering

No assumptions have been made at the point regarding the method used to estimate the target state $P(x)$. For problems that can be accurately represented using linear Gaussian models, the various forms of the Kalman Filter provide efficient estimation and greatly reduce the computation and communication requirements of the algorithms discussed earlier. However, many interesting problems are intrinsically unsuited to this method. For example, the prior distribution of $P(x)$ for a search problem with little prior information will be approximately uniform, and as the search progresses, it may become multi-modal and cannot be usefully represented as Gaussian. A target motion model that is spatially constrained, such as an urban search confined to a street grid, also results in distributions that cannot be approximated as Gaussian.

Recursive Bayesian filtering allows general prior, motion and sensor models, and reduces to the Kalman Filter for the linear Gaussian case. For a series of observations $z_1, \ldots, z_k$ of state x, Bayes rule gives a method to incorporate the final observation z_k, as in Eqn. (27).

$$P(x_k|z_1,\ldots,z_k) = \frac{P(z_k|x_k)P(x_k|z_1,\ldots,z_{k-1})}{P(z_k)} \tag{27}$$

If x is not constant, the subscript represents time and each observation z_i is conditioned on x_i. A prediction step in the form of Eqn. (28) is required to produce $P(x_k|z_1,\ldots,z_{k-1})$ (the prior estimate at time k) from $P(x_{k-1}|z_1,\ldots,z_{k-1})$ (the posterior at time $k-1$). Equations (27) and (28) correspond to the update and prediction steps of the Kalman Filter for the linear Gaussian case. The sensor model is represented by $P(z_k|x_k)$ and the target motion model is $P(x_{k+1}|x_k)$. Recursive Bayesian filtering is used in this example because the bearing-only sensor with restricted range is non-linear and leads to a non-Gaussian state estimate.

$$P(x_k|z_1,\ldots,z_{k-1}) = \int_{x_{k-1}} P(x_{k-1}|z_1,\ldots,z_{k-1})P(x_k|x_{k-1})dx_{k-1} \tag{28}$$

Recursive Bayesian estimation can be decentralized through the exchange of sensor likelihood functions. The observations of the two sensors at time k are independent conditioned on $P(x_k)$ and so the joint likelihood function is simply the product of the individual sensor likelihoods.

$$P(z_k^1, z_k^2|x) = P(z_k^1|x)P(z_k^2|x) \tag{29}$$

Preliminary simulation results

Two bearing-only sensors are simulated tracking a target under individual control and cooperatively estimating its position. Both UAVs maintain identical estimates of the current target position by fusing their observations. Due to the fixed wing flight dynamics, the sensors cannot maintain the slowly moving target in view, and must observe it periodically by circling back to its expected position as shown by the sensor and target trajectories in Fig. 6.

At each time step, the cost function from Eqn. 3 is calculated for $N = 3$: the estimate entropy is predicted three steps into

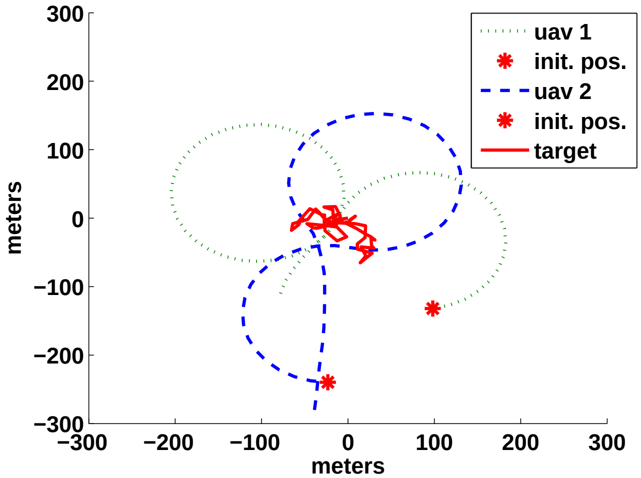

Figure 6. SENSOR AND TARGET TRAJECTORIES FROM SIMULATION. UAV KINEMATIC CONSTRAINTS PREVENT THE SLOWLY-MOVING TARGET FROM BEING CONSTANTLY OBSERVED.

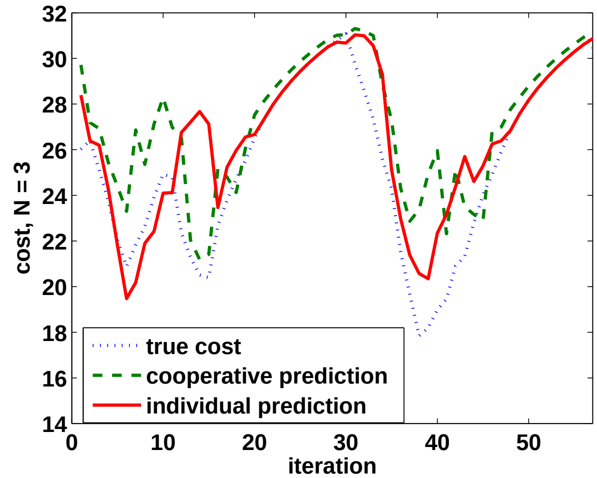

Figure 7. INFORMATION-THEORETIC COST ESTIMATED BY SENSOR 1 WITH HORIZON N = 3. MEAN SQUARE ERROR IS 6.08 AND 4.99 IN THE COOPERATIVE AND SINGLE PREDICTIONS RESPECTIVELY.

the future. After the simulation is complete, the actual receding horizon cost is calculated and compared to the predictions. An individual prediction of J by agent q predicts the entropy conditioned only on observations $[z_k^q, \ldots, z_{k+N}^q]$. A cooperative prediction of J by agent q (cooperating with agent m) predicts the entropy conditioned on observations $[z_k^q, \ldots, z_{k+N}^q, z_k^m, \ldots, z_{k+N}^m]$. The mutual information of the two sensors' observations is estimated using the "most informative limited range" upper bound. Therefore the cooperative prediction is expected to overestimate the entropy when both sensors observe the target.

Figures 7 and 8 show the individually and cooperatively predicted costs from each sensor compared to the actual value from the simulation in figure 6. Averaged over both sensors, the mean square error for the cooperative prediction is 5.44 and the for the individual prediction is 6.99. These correspond to 9.7 percent rms error in the cooperative cost prediction and 11 percent rms error in the individual cost prediction. Although the sensor coupling is overestimated as expected, the cooperative cost prediction is slightly more accurate than the individual one. This improvement will likely be increased with future implementation of a sensor coupling calculation that fully incorporates both sensor models.

These simulation results show that the receding horizon cost function can be calculated based on multiple step predictions of the estimate entropy, in a decentralized manner. The cooperative cost predictions required only communication of planned control and the resulting sensor footprints, rather than system-wide exchange of full sensor models.

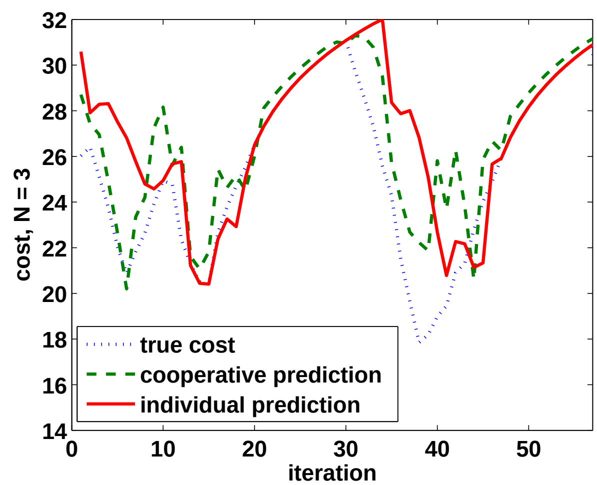

Figure 8. INFORMATION-THEORETIC COST ESTIMATED BY SENSOR 2 WITH HORIZON N = 3. MEAN SQUARE ERROR IS 4.79 AND 8.99 IN THE COOPERATIVE AND SINGLE PREDICTIONS RESPECTIVELY.

CONCLUSIONS AND FUTURE WORK

The contributions of this work are the formulation of cooperative search and track as a receding horizon optimal control problem that fully incorporates probabilistic sensor and target motion models. Searching and tracking behaviors should arise from a single control framework because the transition between the two may not be well-defined for problems with limited sensor range and motion constraints. An information-theoretic cost function is developed to represent a general estimation task for a team of agents, and the structure of the cost function is used to upper bound the coupling between agents. This upper bound determines the required accuracy of the coupling calculation used in optimization.

Future work will begin with implementation of an estimate of the sensor coupling which makes use of the full sensors models, for use when a subset of the observations are tightly coupled. A variety of decentralized optimization methods will be investigated, including both sequential (greedy) and fully cooperative algorithms. The maximum inter-agent coupling may also determine which method is most appropriate. Further simulations including more structured target motion models may show more varied performance and will further display the importance of a combined search and track control goal.

REFERENCES

[1] Gutin, G., and Punnen, A., eds., 2002. *TSP, The Travelling Salesman Problem and its Variations*. Kluwer Academic Publishers.

[2] Ryan, A., Tisdale, J., Godwin, M., Coatta, D., Nguyen, D., Spry, S., Sengupta, R., and Hedrick, J. K., 2007. "Decentralized control of unmanned aerial vehicle sensing missions". In Proceedings of the American Controls Conference.

[3] Stone, L. D., 1989. "What's happened in search theory since the 1975 lanchester prize?". *Operations Research,* **37**(3), May, pp. 501–506.

[4] Spry, S. C., Girard, A. R., and Hedrick, J. K., 2005. "Convoy protection using multiple unmanned air vehicles: Organization and coordination". In Proceedings of the American Control Conference.

[5] Webb, S., and Furukawa, T., 2006. "Belief driven manipulator control for integrated searching and tracking". In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems.

[6] Grocholsky, B., 2005. "Information driven coordinated air-ground proactive sensing". In Proceedings of the IEEE international conference on robotics and automation.

[7] Wong, E., Bourgault, F., and Furukawa, T., 2005. "Multi-vehicle bayesian search for multiple lost targets". In Proceedings of The 2005 IEEE Int. Conf. on Robotics and Automation.

[8] Hoffmann, G. M., Waslander, S. L., and Tomlin, C. J., 2006. "Distributed cooperative search using information-theoretic costs for particle filtes, with quadrotor applications". In Proceedings of the AIAA Guidance Navigation and Control Conference and Exhibit.

[9] Cole, D. T., Goktogan, A. H., and Sukkarieh, S., 2006. "The demonstration of a cooperative control architecture for UAV teams". In Proceedings of the 10th International Symposium on Experimental Robotics.

[10] Grocholsky, B., 2002. "Information-theoretic control of multiple sensor platforms". PhD thesis, The University of Sydney.

[11] Mayne, D., and Michalska, H., 1990. "Receding horizon control of nonlinear systems". *IEEE Transatcions on Automatic Control*, July.

[12] Bertsekas, D. P., Nedic, A., and Ozdaglar, A. E., 2003. *Convex Analysis and Optimization*. Athena Scientific.

[13] Bertsekas, D. P., and Tsitsiklis, J. N., 1997. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, Belmont, MA.

[14] Mathews, G., and Durrant-Whyte, H., 2007. "Decentralised optimal control for reconnaissance". In Proceedings of the Conference on Information, Decision and Control.

[15] Frew, E., and Sengupta, R., 2004. "Obstacle avoidance with sensor uncertainty for small aircraft". In Proceedings of the IEEE Conference on Decision and Control.

[16] Ryan, A., Xiao, X., Rathinam, S., Tisdale, J., Caveney, D., Sengupta, R., and Hedrick, J., 2006. "A modular software infrastructure for distributed control of collaborating unmanned aerial vehicles". In Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit.