Improving object tracking through distributed exploration of an information map

Izaak D. Neveln,[1] Lauren M. Miller,[2] Malcolm A. MacIver,[1,2] and Todd D. Murphey,[2]

***Abstract*— Tracking moving objects requires tight coordination of sensing and movement, in both biological contexts such as prey pursuit and capture, and in target following by mobile robots. Algorithms for target tracking often use a probabilistic map, or information map, of the domain to guide search. Though it is reasonable to expect that the best approach would be to choose control actions driving the robot toward the maximum of this information map, we show improved performance in simulation by using a simple heuristic incorporating the time history of robot movement into the map. Furthermore, our results indicate that as the distribution of robot positions approaches the distribution of the density of information, the variance of the estimate is decreased and tracking improves. We conclude that control actions based solely on information maximization may underperform in information orientated tasks, such as estimation.**

I. Introduction

Mobile robots (as well as their biological counterparts, animals) face uncertain environments from which some certainty must be derived to remain operational (or alive). Robots and animals both require sensory data in order to make sense of their surroundings and perform tasks such as navigation, object identification, etc. Movement of the sensors is often required to maintain a steady stream of useful information about the environment. With movement comes a cost [1], however, so movement strategies must be optimized to gain the most useful information about the environment, leading to an entire field of study called active search.

In robotics, a variety of methods have been proposed in active search. Often, robots gather sensor data to estimate some unknown parameter about their environment. This unknown parameter could be the location of an object to track or a distinguishing feature of an object that requires identification. A perfect sensor without noise would be able to determine this parameter with a single measurement (i.e. there is a one-to-one correspondence between the sensor reading and the value of the parameter). However, most sensors receive noisy measurements, and even in noise-free cases, a measurement could result in multiple possibilities for the value of the parameter. A suitable method for active search should therefore choose a control which will best minimize the variance of the parameter belief function while also resulting in an accurate estimate.

This work was supported by NSF grants IOB-0517683, CMMI-0941674, IIS 1018167, an Office of Naval Research Small Business Technology Transfer grant N00014-09-M-0306 to M.A.M., and National Institute of Health grant T32 HD007418.

[1]Department of Biomedical Engineering. [2]Department of Mechanical Engineering. [1] and [2] at R.R. McCormick School of Engineering and Applied Science at Northwestern University, Evanston, Illinois, USA

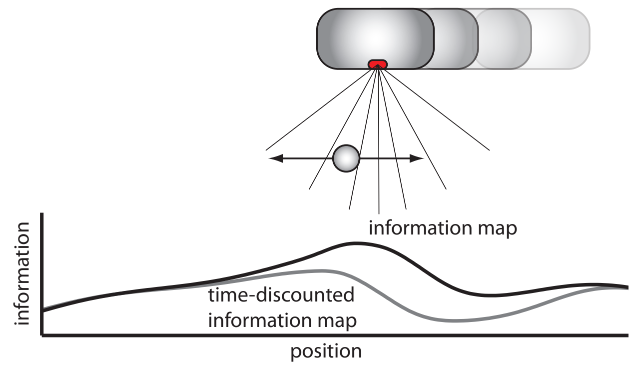

Fig. 1: Schematic of our object tracking method using a time-discounted information map. The robotic sensor is tasked with locating and tracking an object moving in 1D. A simple control strategy might be to ascend the map in order to reach locations of maximum information. Our method discounts the information map according to the recent position history of the robot, allowing for a more distributed exploration of the original map using a simple and real-time control law. This time-discounted information map method results in better estimates over information maximization.

The inherent problem in active search is that the utility of a future measurement must somehow be predicted, given that the measurement itself is dependent on the uncertain parameter. Predicting measurement utility can be accomplished using entropy related metrics [2], [3] or information measures [4]–[6]. As searching the entire space of control actions can be expensive, methods of locally maximizing or approximately maximizing such metrics are often used. For example, local control actions can be chosen in order to maximize some measure on expected measurement utility over a set of candidate control actions [3], [7], [8]. Alternatively, local gradient-based methods based on an information metric [6], [9] can be used to drive sensors towards informative sensor states. While these methods prove sufficient in many situations, local information maximization (info-max for short) methods may fail when the belief, and therefore the expected measurements, are highly uncertain or multimodal. These issues are expected if the parameter is time-varying.

An alternative approach would be using the global maximum of an information map over the whole search domain to to control decisions [10]. One might expect that moving the sensor towards the peak of this information map will result in robust estimation. This approach is however still likely to fail when the belief is uncertain, incorrect, or multimodal, which will be demonstrated in Section III. We show that choosing control actions that are based on the distribution of information, as opposed to the local information or maximum information, can improve performance in these situations.

Previously, we developed a method that calculates a spatial map of expected information using the Fisher Information of an *a priori* measurement model [11]. Using this information map, control actions were chosen over a finite time horizon to generate an ergodic trajectory, meaning the sensor visited locations in the search space for times proportionate to the expected information density of those areas [12]. This ergodic search method was more successful and faster at localizing objects in the presence of distractor objects than other locally greedy methods described above [11].

An ergodic metric for estimation involves analyzing the spatial distribution of measurements over a finite time horizon. Silverman et. al. [11] developed a method of generating forward trajectories that produces ergodic behavior with respect to the current expected information density. Trajectory optimization can however be computationally expensive for dynamic systems and difficult to implement in real time. The benefit of generating a longer trajectory is also less clear when the parameter being estimated is time variant.

Rather than generating an ergodic trajectory forward in time, Mathew and Mezic (2011) [12] calculate an ergodic feedback control law based on past position history. This can still be computationally expensive as it is necessary to calculate the ergodic metric, and extension to moving target tracking is not obvious.

In this work, we develop a method of taking past sensor positions into account when calculating the expected information density map itself. The actual control decision therefore does not involve calculating the potentially expensive ergodic metric, while nevertheless leading to approximately ergodic solutions. We show in simulation that by discounting areas of the information map recently visited, a simple control law can be used to accurately track objects in real time. Using this time-discounted information map results in trajectories that are more distributed (ergodic) over the map, which correlates strongly with a decrease in variance of the estimate. The same control strategy using the original map follows the maximum expected information, but often results in unstable estimates of the moving object. The approximately ergodic (called quasi-ergodic in this paper) approach will likely result in more energetically costly solutions due to increased movement, though information gained could offset this cost [1].

While the control strategies for animals performing active search is a large open area of research, there is some evidence that animals perform costly movements as a trade-off for gaining information. Electric fish swim at a drag-inducing pitched angle to sweep more area with their limited-range sensors [1]. Similarly, blue crabs orient themselves at a drag inducing angle while moving, likely to obtain a better estimate on the local gradient of odorant molecules [13]. Bats orient their ultrasonic clicks offset from prey where Fisher Information is maximized [14]. A recent study in electric fish show a large increase in whole-body oscillations when tracking a moving refuge by means of their electrosensory system compared to when using their visual system [15]. Could these strategies be considering areas of increased information density when planning control? Would these strategies benefit from an ergodic search approach compared to going to areas of maximum information? Investigation into solving these problems in robotics could lead to insight on these questions in animals.

II. METHODS

A. Algorithm overview

Algorithm 1

```
Init. d(0), V_0(0), Υ(θ, d), ϵ
Init. p(θ) to a uniform distribution
Calculate the Fisher Information I(θ, d) (Eqn. 1)
while True do
    Calculate expected information density (EID) map using p(θ) and I(θ, d) (Eqn. 2)
    Calculate sensor position history to discount EID (Fig. 2)
    Update control f from time-discounted EID (tEID) (Eqn. 5)
    Take measurement and update p_i(θ) (Eqn. 3)
    i=i+1
end while
```

The algorithm presented in this paper reformulates the algorithm in Silverman et. al. (2013) [11]—which was designed to localize a static object—in order to track a moving object. Previously, a trajectory was planned for a finite block of time, executed, and then repeated with the updated belief of the object location (more generally, object parameter θ). The current algorithm takes into account the recent trajectory of the sensor along with the expected information density (EID) of the parameter when executing the subsequent control. For example, the sensor will explore areas of lower EID if areas of higher EID had been explored recently. An overview of this method is shown in Algorithm 1. As in Silverman et. al. (2013) [11], we assume knowledge of the measurement model for θ, given by $v = \Upsilon(\theta, d) + \delta$. Υ is a function of both the deterministic sensor position d as well as the unknown parameter θ and can vary depending on the type of sensor used and the parameter being estimated. δ adds a zero mean noise with variance σ^2. The measurement model could be derived from first principles if the physics of the sensor are well known, or empirically obtained through experiment.

B. Fisher Information

Given the measurement model v, we can calculate the Fisher Information (FI) as in Silverman et. al. (2013) [11]. FI correlates to the amount of information a measurement will provide at sensor location x for a specific value of θ. Assuming Gaussian noise of the measurement, FI can be calculated as

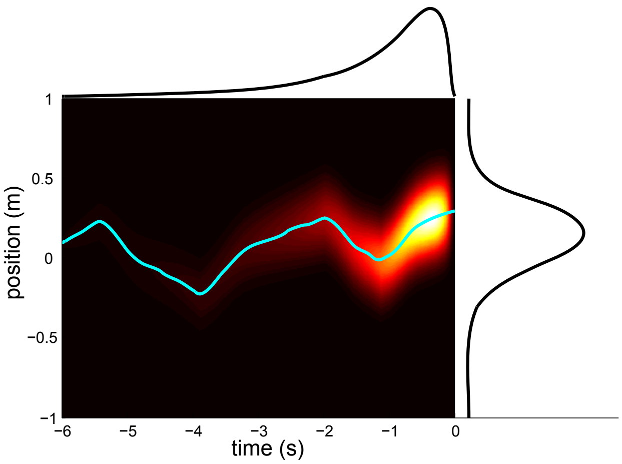

Fig. 2: Calculation of position history function for discounting the expected information density. The blue line shows a sample trajectory of the sensor over the last 6 seconds, where time = 0 indicates the current time. The colormap indicates how the function is created over time. The top plot indicates the level of the discounting as a function of time. Bright yellow indicates areas of high magnitude for the discounting function. The right plot is the sum of the discounting function over all time, which is then used with the EID so that positions of high magnitude (areas recently visited by the sensor) are discounted heavily whereas low magnitudes (areas not recently visited by the sensor) are not discounted.

$$\mathcal{I}(\theta, d) = \int_v \left(\frac{\partial p(v|\theta)}{\partial \theta}\right)^2 \frac{1}{p(v|\theta)} dv. \tag{1}$$

The belief about the value of θ is represented by the PDF $p(\theta)$ (also called the belief) and evolves as measurements are collected (see Section II-D). To calculate the expected information density (EID), we take the expectation of the Fisher Information over the belief of the parameter using

$$\Phi(x) = \int_\theta \mathcal{I}(\theta, x) p(\theta)\, d\theta. \tag{2}$$

The calculation of FI ($\mathcal{I}$) and EID (Φ) are unchanged from the algorithm presented in Silverman et. al. (2013) [11].

C. Sensor position history and control update

Previously, in Silverman et. al. (2013) [11], an ergodic sensor trajectory was calculated for a finite time horizon using optimal control methods. In that case, a trajectory was ergodic if the distribution of locations that the trajectory visits matches the statistics of the EID. Therefore, the sensor would spend more time in areas of high EID, and proportionately less time in areas of lower EID.

When the unknown parameter has the possibility to evolve over time, such as when tracking a moving object, calculating an ergodic trajectory over a finite time horizon becomes impractical because of how rapidly the EID is changing.

Specifically, since an ergodic trajectory requires matching the sampling of a space to the statistics of the EID, then if the EID is changing, there is no assurance that this matching will occur. To accommodate this, we introduce two forms of "forgetting" into our algorithm, so that guidance by memory (the information map) is less prone to error. The first form is to forget, with a certain time constant, the EID of recently visited locations. We call the result the time-discounted EID, or tEID. A schematic of this method is shown in Fig. 2. The second is to forget the current estimate of the parameter, which will be further explained below. For this work, the time-constants associated with these two forms of forgetting were tuned manually by visualizing the resulting trajectories, whereas future versions of this algorithm would benefit from automatic tuning based on the frequency spectrum of the estimated parameter.

A control update was then calculated using the tEID. The specific control law to use may depend on the dynamics of the robot and the dimenionality of the problem. We chose to focus on a very simple control law to show the effectiveness of using a time-discounted map, which is given in Section II-E.

D. Bayesian update and evolution of the belief

The PDF of the parameter $p(\theta)$ is updated through a similar Bayesian update as in Silverman et. al (2013) [11]. The update is given by

$$p(\theta|\boldsymbol{V}) = \eta p(\boldsymbol{V}|\theta) p(\theta), \tag{3}$$

where $p(\theta)$ is the prior belief, $\boldsymbol{V}$ is the measurement, $p(\boldsymbol{V}|\theta)$ is the innovation, and η is a normalization factor. For estimating an unchanging parameter, the innovation is calculated from

$$p(\boldsymbol{V})|\boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[\frac{(\Upsilon(\boldsymbol{\theta}, d) - V)^2}{\sigma^2}\right]. \tag{4}$$

However, here we model the dynamics of the parameter (position in the case of a moving object) as a Gaussian distribution centered at the current value of the parameter (also known as a stochastic maximum velocity method). As the parameter is represented by a probability distribution, the calculation amounts to convolving the distribution for the parameter with the Gaussian distribution representing the possible evolution of the parameter. Therefore, the belief of the parameter flattens out over time if no new measurements are taken, effectively 'forgetting' some knowledge of the location of the object. This flattened distribution then becomes the prior in Eqn. 3. This procedure allows the estimate of the parameter to change over time, even if nothing is known deterministically about how the parameter might evolve.

E. Example Problem

We simulated a single sensor that gathers some reading as it passes by an object in 1D. The global position of that object is the unknown parameter θ. $\Upsilon(\theta, x)$ models the sensor reading as a function and can vary depending on the type of sensor or the object. We chose to compare three sensor models as well as two control strategies. The first strategy

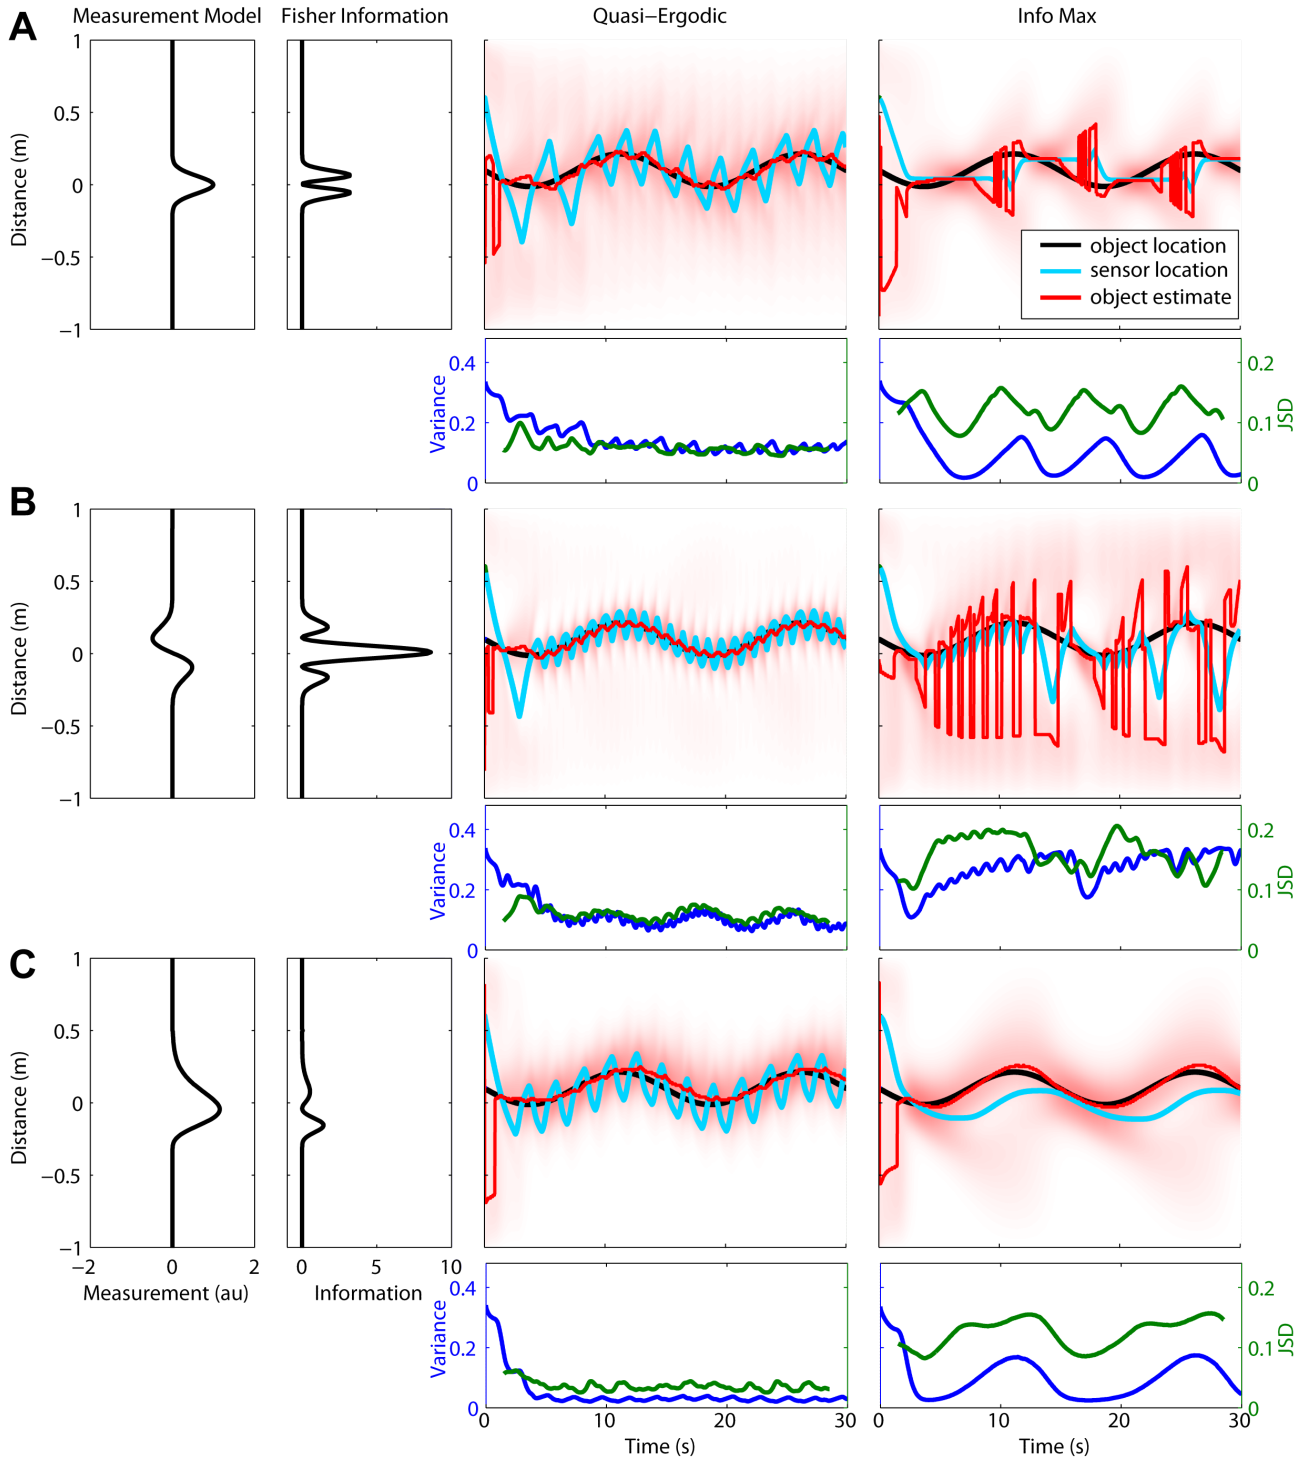

Fig. 3: Quasi-ergodic control vs. info-max control for tracking an object. The left column shows the different measurement models tested. The first column shows the sensor reading as an object passes by a sensor located at 0. The sensor models would shift up and down according to the location of the sensor. A monophasic, symmetric measurement model is used in (A), a biphasic, symmetric model is used in (B), and a monophasic asymmetric measurement model is used in (C). The Fisher Information for each measurement model is shown in the second column. High FI corresponds to the locations where the magnitude of the spatial derivate of the measurement model is also high. The third and fourth columns show simulated trials as the sensor attempts to track an object using either the quasi-ergodic control method (third column) or the info-max method (fourth column). The black trace indicates the trajectory of the object and the cyan trace indicates the trajectory of the sensor. The shaded red shows the PDF of the estimated object position over space and time, where the red trace indicated the estimate of the object position based on the location of the maximum value of the PDF. In all cases, the period of the object oscillation in 15 seconds and two full periods are shown. The plots underneath each trial detail the time evolution of the variance of the estimate (blue) as well as the Jensen-Shannon divergence measure of ergodicity (green). The signal to noise ratio was 25 dB in all cases.

takes into account the time-discounted EID as described above. The second strategy is to follow the maximum of the EID, irrespective of where the sensor has been. We allow the position of the object to oscillate at various frequencies,

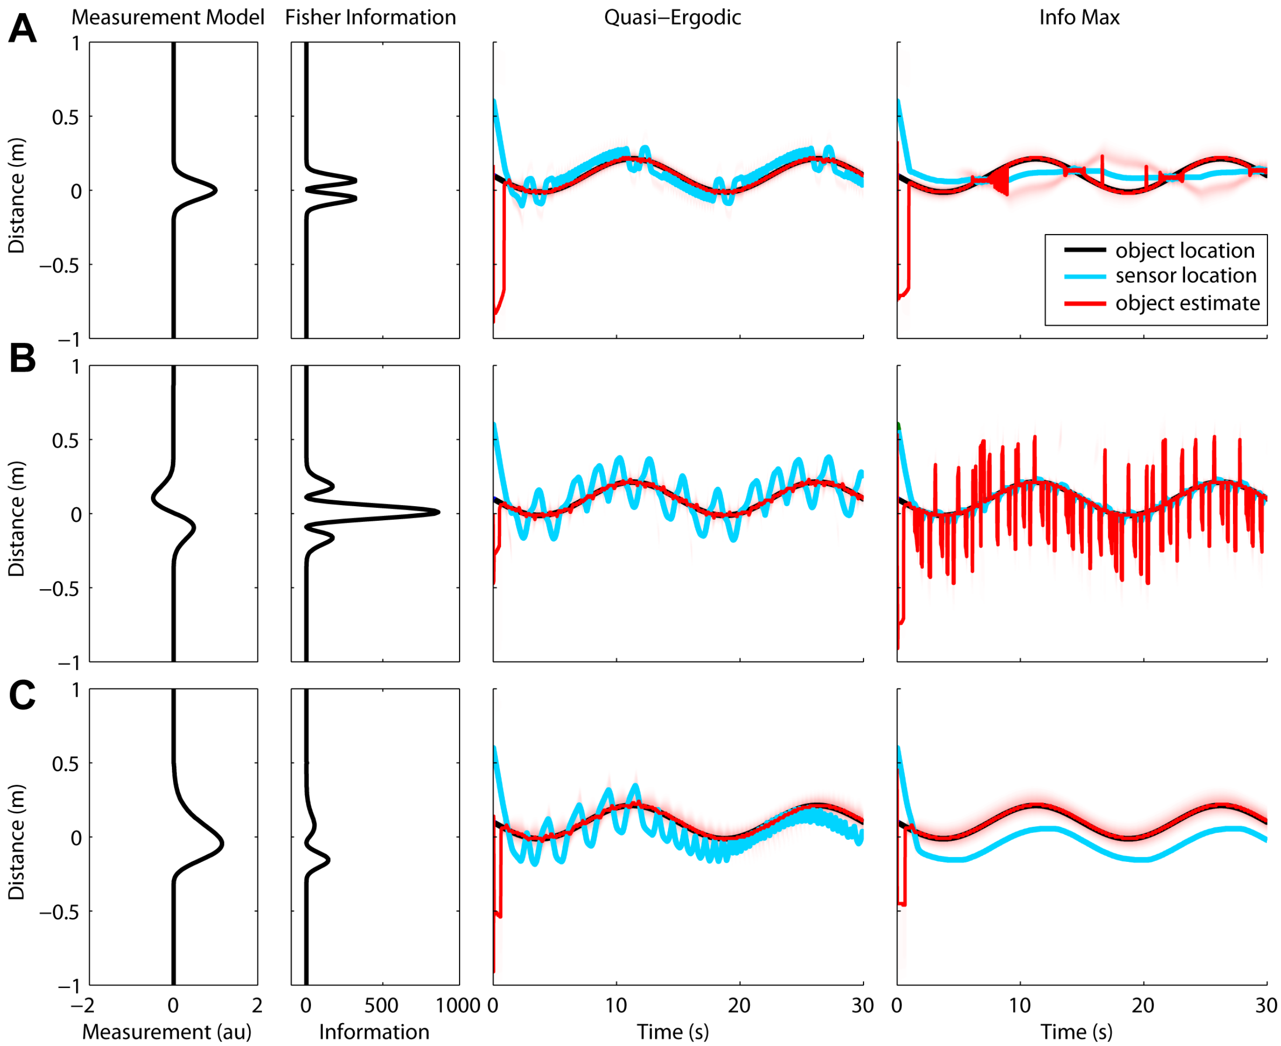

Fig. 4: Quasi-ergodic control vs. info-max control for tracking an object with lower noise. These trials are identical to those in Fig. 3 except that the amplitude of the noise modeled was decreased by one order of magnitude. The signal to noise ratio was 45 dB in all cases. See Fig. 3 for more details.

keeping the amplitude of the velocity profile constant.

We chose a very simple control law based on the time-discounted EID, where the input f drives the sensor toward the peak of the tEID with a magnitude inversely proportional to the tEID at the current location, given by

$$f \propto \frac{\text{sign}[(\text{argmax}_x \, \text{tEID}(x)) - d]}{\text{tEID}(d)} \quad (5)$$

where tEID is the time-discounted EID function over the domain x and d is the current position of the sensor. By using the original EID rather than the time-discounted EID, the result is to move towards the global maximum of the information map. The sensor dynamics consists of a simple second order linear system with an inertial term, a damping term, and no stiffness with full actuation in 1D given by

$$m\ddot{d} + b\dot{d} = f, \quad (6)$$

where m and b are the mass and damping coefficient of the sensor. This model represents the dynamics of our underwater electrosensory robot with which we will perform future experiments [16].

F. Performance Measures

Three performance measures were used to quantify the results. First, the norm was calculated for the error signal between the estimated and actual object position during the steady state phase of tracking (the last 20 seconds). Two error signals were derived based on the two possible ways to estimate object position, either by the mean location of the belief function, or the maximum location of the belief function. These norms are reported in Table I.

Second, the variance of the belief over the time course of the trajectory was calculated. This measure relates the confidence of the estimate over time. Third, the Jensen-Shannon divergence (JSD) was used to measure the ergodicity of the sensor trajectory related to the information map (the EID). In short, the JSD is a measure of distance between two distributions. We calculated the distribution of sensor positions as well as the average EID in a moving window of 3 seconds. Therefore, a value of JSD close to zero indicates that the position history of the sensor matches the EID, and therefore represents high ergodicity. Both the variance of the belief and the JSD are plotted for the trials shown in Fig. 3.

III. RESULTS

We present trials where the sensor must track a moving object in a high noise environment (SNR = 25 dB), comparing the quasi-ergodic approach to the information maximizing approach using a variety of measurement models. Fig. 3A shows results when using a symmetric, monophasic measurement model, common in many types of sensors for simple objects such as sonar detecting a prey. Fig. 3A shows that the estimate of the position converges quickly and remains close to the actual position for the quasi-ergodic method, while the info-max method rapidly flips its estimate between two possible positions of the object. These rapid changes in the estimate correspond to times when the variance of the belief is increasing and the current ergodicity is minimized (JSD is increased). Norms on the tracking error are reported in Table I.

Fig. 3B shows a similar trial for a biphasic measurement model, similar to the one used in [11], as it describes the voltage reading of our active electrosensing system as it passes by an object. The Fisher Information for this type of measurement model has a large peak centered with the object with two smaller peaks on either side. Here again, the quasi-ergodic method tracks the object reliably, whereas the info-max method maintains a belief function with high variance of the object position, where the estimate jumps radically.

Last, Fig. 3C show another monophasic measurement model, only now it is asymmetric, resulting in one side of the object to contain more Fisher Information than the other side. The sensor position is biased towards the side with more Fisher Information as expected for both the quasi-ergodic case and the info-max case. While both methods track the object reasonably well, the variance of the estimate is unstable and generally greater for the info-max method. Also, when the object changes direction, the estimate for the info-max method briefly deviates from the sinusoidal pattern of the object.

We also simulated the trials with lower noise to determine if the SNR plays a role in how each method performs. The results for tracking an object moving sinusoidally with a period of 15 seconds and SNR $= 45$ dB is shown in Fig. 4. Norms of the tracking error are reported in Table I.

While not shown in this article, we also simulated objects oscillating with shorter periods, such as 10 and 5 seconds, but keeping the amplitude of the velocity profile fixed. The quasi-ergodic and info-max methods exhibited similar behavior for these higher frequency tracking simulations. We also varied the initial position of the sensor and did not see any effect on the tracking behavior.

IV. DISCUSSION

These results indicate that movement outside of regions of maximum expected information is often necessary to maintain good estimates while tracking objects, especially in high-noise environments. Also, the quasi-ergodic method shown here is more robust to changes in the measurement model as well as noise.

Measurement Model	Control Method	SNR = 25 dB Max	SNR = 25 dB Mean	SNR = 45 dB Max	SNR = 45 dB Mean
Monophasic symmetrical	Quasi-ergodic	0.71	2.20	0.14	0.46
	Info-max	4.68	3.81	0.51	2.32
Biphasic symmetrical	Quasi-ergodic	0.60	2.73	0.24	0.61
	Info-max	14.5	4.48	4.69	4.82
Monophasic asymmetrical	Quasi-ergodic	1.27	1.79	0.23	0.53
	Info-max	0.88	2.80	0.19	0.34

TABLE I: Norm of the tracking error over the last 20 seconds of each trial from Figs. 3 and 4. The error was calculated as the difference between the actual object position and the estimate of the object position. The estimate could either be calculated as the mean of the belief function, or the maximum of the belief function. The values for the norm have units of meters. The heigth of the bars are proportional to the values of the norms, except for the entry of 14.5 which is about 3 times as large as any other value so is.

The measurement model in Fig. 3A represents a simple sensor and a simple object. However, the symmetry of the measurement model poses problems. First, there is no unique maximum in the Fisher Information, as two equal peaks are offset from the center of the object. If the info-max method is used to track an object with this measurement model, the PDF becomes bimodal, as a measurement from one side of the object could indicate two possible locations for the object. Therefore, an estimate based on the mean of the PDF would average these two modes, but an estimate based on the maximum of the PDF might switch rapidly between the two modes as it does for the info-max method. Indeed, the norm of the tracking error is relatively high for both estimates. The quasi-ergodic method, since it receives measurements from both sides of the object, maintains a unimodal PDF; therefore the estimate is unambiguous and the variance remains constant. Even for low-noise situations as shown in Fig. 4A, the initial sweeps of the sensor allow the quasi-ergodic method to quickly settle on a unimodal PDF, while the bimodal PDF persists for the info-max method. This measurement model is likely similar to that of a bat detecting a small prey, which also distributes its ultrasonic clicks on either side of the prey [14].

The measurement model in Fig. 3B is similar to the model of measurements from our artificial electrosensors as an object passes by them. This measurement model is interesting because the highest Fisher Information is right at the center of the object, but the actual measurements at that location are very similar to those that are measured far away from the object. Therefore, once the PDF focuses on an estimate of the object, oscillations of the sensor allow it to disambiguate between the object being centered or being far away. The info-max method has difficulty with this disambiguation, persisting even in low noise situations (Fig. 4B).

Both methods are able to track the asymmetric measurement model shown in Fig.3C and the norm for the tracking error using the maximum as the estimate is lower for info-max vs. quasi-ergodic. However, the info-max method exhibits a bimodal PDF similar to that from the symmetric measurement model in high noise situation resulting in large fluctuations in the variance and a poorer norm when the mean of the belief is used as the estimate. In low noise situations, the quasi-ergodic method and info-max method both perform well (the norms on the tracking error are actually slightly lower for info-max), and the trajectory of the quasi-ergodic method appears to converge on the info-max trajectory.

An interesting result which is especially apparent in the case of the two monophasic measurement models is the correlation between the variance of the estimate and the ergodicity as measured by the JSD. The JSD varies directly with the variance, indicating that times where the ergodicity is high (low JSD), the estimate actually improves. For the info-max control cases, there are times where the trajectory happens to be more ergodic as a result of the object movement, and it is at these times when the estimate improves. The quasi-ergodic method, which simply uses the time-discounted information map (tEID), achieves the goal of maintaining higher and more stable ergodicity. These results show that in many object tracking situations, even if a good estimate is quickly obtained, it can be detrimental to try to lock the movement of the sensor to the movement of the object.

V. Future work

Sensor oscillation is a common phenomenon in biological systems, such as full body oscillations of electric fish in tracking behaviors [15], or small amplitude oscillation in eyes to avoid adaptation of retinal cells. Future work will involve testing models similar to the quasi-ergodic method of tracking with behavioral data of animals performing active sensing. Also, we plan to implement these methods on robotic systems to test their efficacy on real sensors sensing real object with natural levels of noise. We plan to adapt our methods to work for estimating more parameters with higher dimensional movement, as well automatically tune the time constants according to the frequency spectrum of the estimate. Also, we would like to incorporate energy models in which we impose a cost on movement allowing us to optimize the trajectories of sensors to gain the maximum amount energy (gained in the form of information, lost in the form of movement and sensing costs).

References

[1] M. A. MacIver, A. A. Shirgaonkar, and N. A. Patankar, "Energy-information trade-offs between movement and sensing," *ploscb*, vol. 6, 2010.

[2] C. Leung, S. Huang, N. Kwok, and G. Dissanayake, "Planning under uncertainty using model predictive control for information gathering," *Robotics and Autonomous Systems*, 2006.

[3] C. Kreucher, K. Kastella, and A. O. Hero, "Sensor management using an active sensing approach," *Signal Processing*, vol. 85, no. 3, pp. 607–624, 2005.

[4] G. Zhang and S. Ferrari, "An adaptive artificial potential function approach for geometric sensing," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*. IEEE, 2009, pp. 7903–7910.

[5] J. Tisdale, Z. Kim, and J. K. Hedrick, "Autonomous UAV path planning and estimation," *IEEE Robotics and Automation Magazine*, pp. 35–42, June 2009.

[6] B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, "Cooperative air and ground surveillance," *Robotics Automation Magazine, IEEE*, vol. 13, no. 3, pp. 16–25, 2006.

[7] J. R. Solberg, K. M. Lynch, and M. A. MacIver, "Active electrolocation for underwater target localization," *International Journal of Robotics Research*, vol. 27, no. 5, pp. 529–548, 2008.

[8] H. J. S. Feder, J. J. Leonard, and C. M. Smith, "Adaptive mobile robot navigation and mapping," *International Journal of Robotics Research*, vol. 18, no. 7, pp. 650–668, 1999.

[9] C. Kreucher, J. Wegrzyn, M. Beauvais, and R. Conti, "Multiplatform information-based sensor management: an inverted UAV demonstration," in *Defense and Security Symposium*. International Society for Optics and Photonics, 2007, pp. 65 780Y–65 780Y.

[10] J. Toh and S. Sukkarieh, "A Bayesian formulation for the prioritized search of moving objects," in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, 2006, pp. 219–224.

[11] Y. Silverman, L. Miller, M. A. MacIver, and T. D. Murphey, "Optimal planning for information acquisition," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, 2013, pp. 5974–5980.

[12] G. Mathew and I. Mezic, "Metrics for ergodicity and design of ergodic dynamics for multi-agent system," *Physica D-nonlinear Phenomena*, vol. 240, no. 4-5, pp. 432–442, Feb 2011.

[13] M. J. Weissburg, C. P. James, D. L. Smee, and D. R. Webster, "Fluid mechanics produces conflicting constraints during olfactory navigation of blue crabs, *Callinectes sapidus*," *J. Exp. Biol.*, vol. 206, no. 1, pp. 171–180, 2003.

[14] Y. Yovel, B. Falk, C. F. Moss, and N. Ulanovsky, "Optimal localization by pointing off axis," *Science*, vol. 327, no. 5966, pp. 701–704, 2010.

[15] S. A. Stamper, E. Roth, N. J. Cowan, and E. S. Fortune, "Active sensing via movement shapes spatiotemporal patterns of sensory feedback," *J. Exp. Biol.*, vol. 215, no. Pt 9, pp. 1567–1574, May 2012.

[16] S. Sefati, I. D. Neveln, E. Roth, T. R. Mitchell, J. B. Snyder, M. A. Maciver, E. S. Fortune, and N. J. Cowan, "Mutually opposing forces during locomotion can eliminate the tradeoff between maneuverability and stability," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 110, no. 47, pp. 18 798–18 803, Nov 2013.