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 the position of moving objects requires tight coordination of sensing and movement, in both biological contexts such as prey pursuit and capture, and in target localization by mobile robots. Algorithms for target tracking often use a probabilistic map, or information map, of the domain to guide active 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 under-perform in information orientated tasks, such as the estimation of moving target positions.

I. INTRODUCTION

Mobile robots and animals share the common problem of exposure to uncertain environments from which some certainty must be derived to remain viable. Robots and animals both require sensory data in order to make sense of their surroundings and perform tasks such as navigation, object identification, etc. For short range sensors, movement is often required to maintain a steady stream of useful information about the environment. With movement comes an energetic 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. An ideal 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 action which will
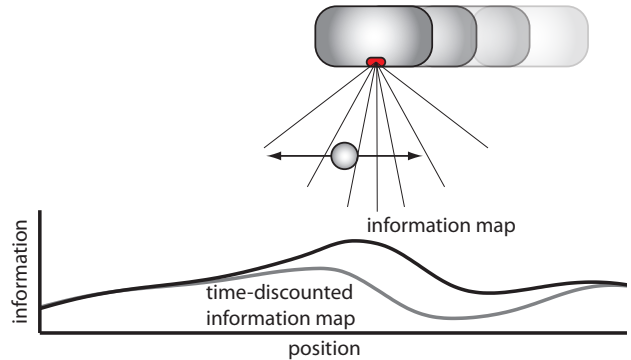
[1]Department of Biomedical Engineering. [2]Department of Mechanical Engineering. [1] and [2] at R.R. McCormick School of Engineering and Applied Science at Northwestern University, Evanston, Illinois, USA

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.

best minimize the variance of the parameter belief function while also resulting in an accurate estimate.

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 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) methods may fail when the belief, and therefore the expected measurements, are highly uncertain or multimodal. These issues are particularly likely to arise if the parameter is time-varying.

Alternatively, the global maximum of an information map over the whole search domain can be used to inform 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. We achieved distributed active sensing guiding our sensor using an information map that has been discounted according to recent sensor position history, allowing us to track the position of a moving target (see Fig. 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 upstream at a drag inducing angle after detecting an odorant, likely to obtain a better estimate on the local gradient of odorant molecules [11]. Bats orient their ultrasonic clicks slightly away from the location of targets in regions where Fisher Information is maximized [12]. 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 [13]. Could these strategies be considering areas of increased information density when planning control? Would these strategies benefit from a more distributed sensing 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

1: Init. $d(0), \boldsymbol{V_0}(0), \Upsilon(\theta, d), \epsilon$
2: Init. $p(\theta)$ to a uniform distribution
3: Calculate the Fisher Information $I(\theta, d)$ (Eqn. 1)
4: **while** True **do**
5:     Calculate expected information density (EID) map using $p(\theta)$ and $I(\theta, d)$ (Eqn. 2)
6:     Calculate sensor position history to discount EID (Fig. 2)
7:     Update control f from time-discounted EID (tEID) (Eqn. 5)
8:     Take measurement and update $p_i(\theta)$ (Section II-D)
9:     i=i+1
10: **end while**

The algorithm presented in this paper reformulates the algorithm in Silverman et. al. (2013) [14]—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) [14], 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.

It might be natural to choose control strategies that produce trajectories maximize information by moving toward the maximum point in the EID. In Algorithm 1, this information maximizing strategy can be accomplished by omitting step 6 and by using the EID in step 7. The main result in this paper is that the information maximizing strategy fails dramatically even in the simplest 1D tracking tasks; distributing information acquisition produces more robust estimations of the parameter being tracked, at least in these 1D tasks.

B. Fisher Information

Given the measurement model v, we can calculate the Fisher Information (FI) as in Silverman et. al. (2013) [14]. 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

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

The belief about the value of θ is represented by the probability density function (PDF) $p(\theta)$ 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) [14].

C. Sensor position history and control update

To accommodate the dynamic nature of the object and the associated EID, we introduce two forms of "forgetting" into our algorithm, so that guidance by memory (our belief and corresponding information map based on previous measurements) 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
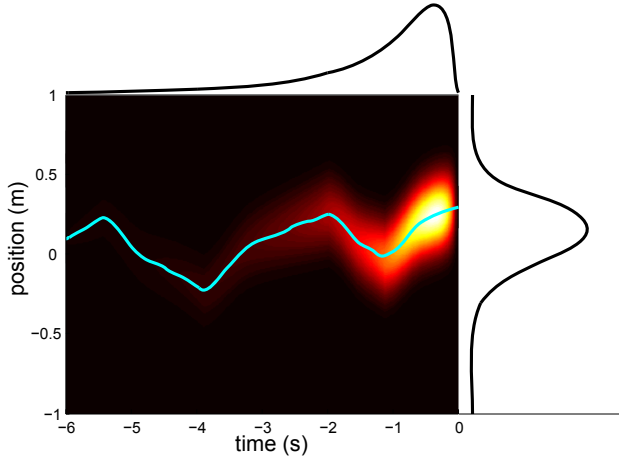
Fig. 2: Calculation of position history function for discounting the expected information density. The blue line shows an example trajectory of the sensor over the last 6 seconds, where time = 0 indicates the current time. The colormap indicates how the function is created over time. The top plot indicates 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 combined 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.

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 dimensionality 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. For the rest of this article, this method of control is labeled the tEID method.

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) [14]. The update is given by

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

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

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

However, here we model the evolution 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 at the previous iteration 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. This update method, chosen for simplicity, could easily be replaced by other methods, such as a particle filter.

E. Example Problem

We simulated a single sensor that gathers measurements as it passes by an object in 1D. The global position of that object is the unknown parameter θ. $\Upsilon(\theta, d)$ models the sensor measurement as a function of sensor position and the parameter. We chose to compare three sensor models as well as two control strategies. The first strategy takes into account the tEID as described above. The second strategy follows the maximum of the EID, irrespective of where the sensor has been. We analyzed the results for trials where the object oscillates with a period of 15 seconds and an amplitude of 0.11 meters. We experimented with other object trajectories though results were qualitatively similar.

We chose a simple control law based on the tEID, 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

$$f \propto \frac{\operatorname{argmax}_x \text{tEID}(x) - d}{\text{tEID}(d)|| \operatorname{argmax}_x \text{tEID}(x)) - d||_2} \qquad (5)$$

where tEID is the time-discounted EID function over the domain x and d is the current position of the sensor. The info-max method calculates the input f in the same way, the only difference being that the EID is used instead of the tEID. The result is that the sensor moves towards the global maximum of the information map. The sensor dynamics consist of a simple second order linear system with an inertial term, a damping term, and no stiffness with full actuation in 1D given by

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

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 [15].

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
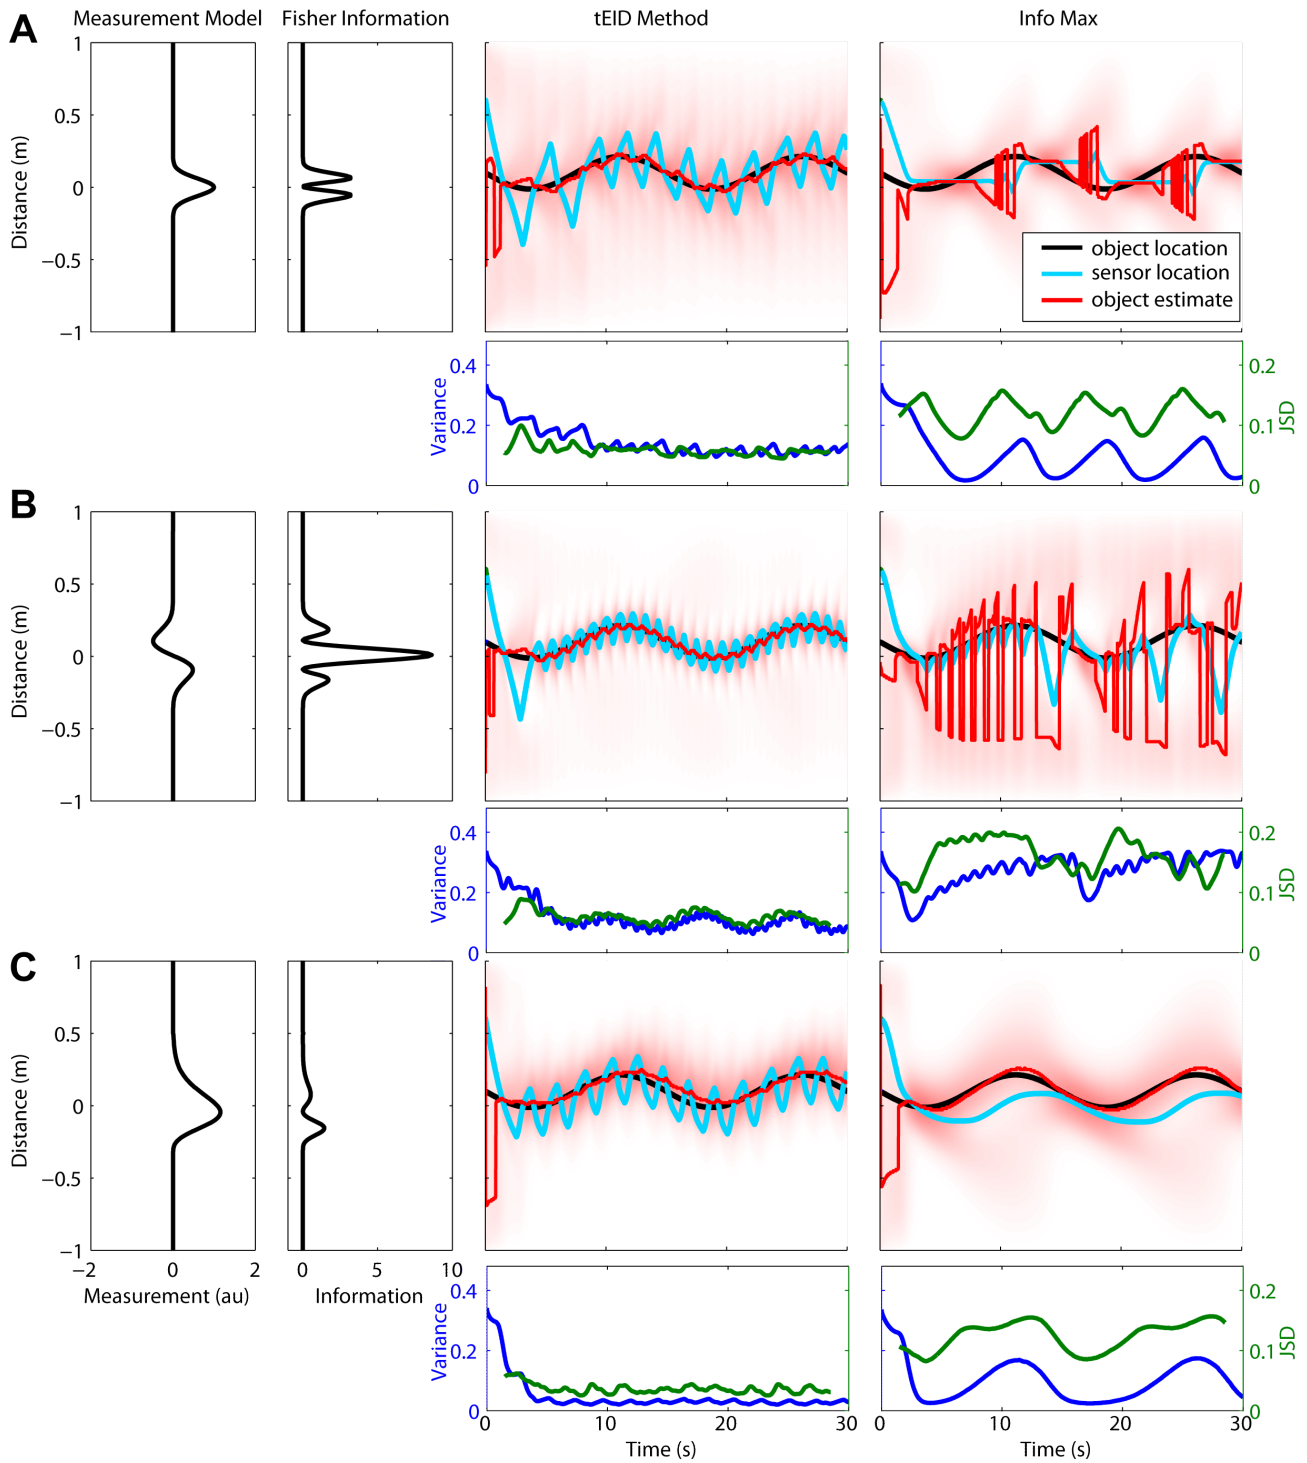
Fig. 3: tEID method vs. info-max control for tracking an object. The left column shows the different measurement models tested. The first column shows the sensor reading as an object passes by a sensor located at $d = 0$. The sensor models would shift up and down according to the location of the sensor. 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 with respect to the parameter is also high. The third and fourth columns show simulated trials as the sensor attempts to track an object using either the tEID 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 is 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.

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.

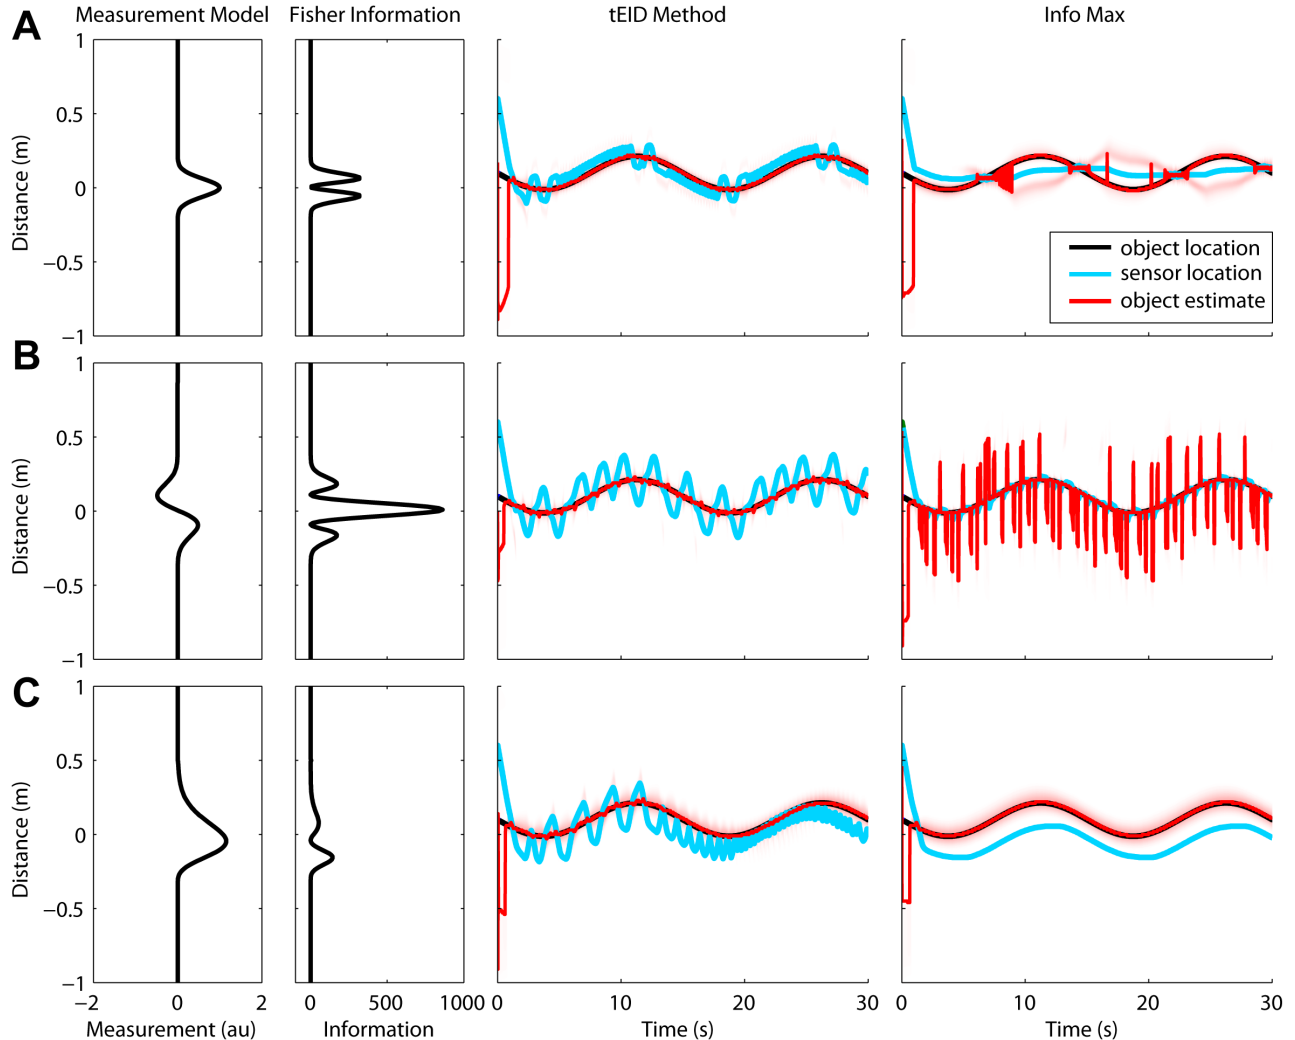Second, the variance of the belief over the time course

Fig. 4: tEID method 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.

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 spatial statistics of the position history of the sensor matches the EID (i.e. the spatial statistics are nearly ergodic with respect to the EID). The motivation for assessing performance using the JSD is based on previous experimental work using ergodicity directly to generate trajectories for target localization [14]. We sought to assess whether or not the heuristic control law presented here produces trajectories that are closer to ergodic than the info-max approach, and whether or not the ergodicity corresponds to better performance of the estimation, as was shown in the target localization task [14]. Both the variance of the belief and the JSD are plotted for the trials shown in Fig. 3.

III. RESULTS

We simulated trials where a mobile sensor must estimate the location of a moving object in a high noise environment (signal to noise ratio (SNR) = 25 dB), comparing the tEID approach to the information maximizing approach using a variety of measurement models, $\Upsilon(x, \theta)$. 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 tEID method, while the info-max method rapidly flips its estimate between two possible positions of the object because of a bimodal belief function. These rapid changes in the estimate correspond to times when the variance of the belief and JSD are increasing. 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 Silverman et. al (2013) [14], as it describes the voltage reading of our active electrosensing system as it passes by an object. The FI for this type

Measurement Model	Control Method	SNR = 25 dB		SNR = 45 dB	
		Max	Mean	Max	Mean
Monophasic symmetrical	tEID	0.71	2.20	0.14	0.46
	Info-max	4.68	3.81	0.51	2.32
Biphasic symmetrical	tEID	0.60	2.73	0.24	0.61
	Info-max	14.5	4.48	4.69	4.82
Monophasic asymmetrical	tEID	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 height of the bars is proportional to the values of the norms.

of measurement model has a large peak centered with the object with two smaller peaks on either side. Here again, the tEID 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 containing more FI than the other side. We therefore observe that the sensor trajectory is biased towards the side with more FI as expected for both the tEID case and the info-max case. While both methods track the object reasonably well, the variance of the estimate fluctuates over time and is generally larger on average 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 same three trials with lower noise to determine if the SNR plays a role in how the tEID or info-max methods perform. The results for tracking an object moving sinusoidally with a period of 15 seconds and SNR = 45 dB are 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 5 and 10 seconds, but keeping the amplitude of the velocity profile fixed. The tEID 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 tEID method shown here is more robust to different measurement models as well as differences in noise levels.

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 FI, 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 tEID 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 tEID 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 [12].

The measurement model in Fig. 3B is similar to the model of measurements from artificial electrosensors used in previous studies [14]. This measurement model is interesting because the highest FI is right at the center of the object, but the actual measurements at that location are 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 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. tEID. However, the info-max method exhibits a bimodal PDF similar to that from the symmetric measurement model in high noise situations 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 tEID 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 tEID 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 tEID 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.

The simple control law based on the tEID generated distributed sensing trajectories and works well in the 1D

tracking example. However, a more rigorous control law could be implemented and would likely be necessary for more complicated tasks. For example, one could follow the gradient of an ergodicity metric [16], or optimize a trajectory by maximizing future ergodicity [14]. Both of these examples would have to be adapted to allow for a time-varying parameter and information map, perhaps by using an information map that includes time-history information such as the tEID used in this paper.

V. FUTURE WORK

Sensor oscillation, which is observed in the trajectories exhibited by the tEID method, is a common phenomenon in biological systems, such as full body oscillations of electric fish in tracking behaviors [13], or small amplitude oscillation in eyes to avoid adaptation of retinal cells. Future work will involve testing models similar to the tEID 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 using real sensors taking measurements of real objects with natural levels of noise. We plan to adapt our methods to work for estimating more parameters with higher dimensional movement and to automatically tune the time constants according to the frequency spectrum of the estimate. Finally, we would like to incorporate energy models in which we impose a cost on movement allowing us to optimize the trajectories of sensors to maximize energy that is gained in the form of information but 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," *PLoS Computational Biology*, 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," *The 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] 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*," *Journal of Experimental Biology*, vol. 206, no. 1, pp. 171–180, 2003.

[12] 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.

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

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

[15] 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.

[16] 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.