Spatial Interpolation for Robotic Sampling: Uncertainty with two Models of Variance

Young-Ho Kim, Dylan A. Shell, Colin Ho and Srikanth Saripalli

Abstract Several important forms of robotic environmental monitoring involve estimating a spatial field from comparatively few measurements. A number of researchers use linear least squares estimation techniques, frequently either the geostatistical Kriging framework or a Gaussian Process regression formulation, that provide estimates of quantities of interest at unmeasured locations. These methods enable selection of sample locations (e.g., for adaptive sampling) by quantifying uncertainty across the scalar field. This paper assesses the role of pose uncertainty and measurement error on variance of the estimated spatial field. We do this through a systematic empirical comparison of scalar fields reconstructed from measurements taken with our robot using multiple imperfect sensors and actively estimating its pose. We implement and compare two models of variance: Kriging Variance (KV) and Interpolation Variance (IV), illustrating that the latter —which has not been used in a robotics context before— has several advantages when used for online planning of sampling tasks. Using two separate experimental scenarios, we assess the estimated variance in scalar fields constructed from measurements taken by robots. Physical robots sampling within our office building suggest that using IV to select sampling sites gathers more data for a given time window (45% more than KV), travels a shorter distance to collect the same number of samples (25% less than KV), and has a promising speed-up with multiple robots. Water quality data from an Autonomous Underwater Vehicle survey of Lake Pleasant, AZ. also show that IV produces better qualities for given a distance and time.

Keywords: Environmental monitoring, robotic sampling, Measurement and position error

1 Introduction

Large-scale environmental monitoring is a particularly promising application for robots [8]. Robots have already begun to collect oceanographic data sets of unprecedented scale and resolution (*e.g.*, [4]). The underlying challenge addressed by such systems stems from the fact that the measured data are sparse compared to the large spatial areas/volumes of interest. Mobility makes adaptive strategies for *in situ* sampling possible but leads naturally to the question: *"given the data already captured, where should the robots go in order to sample further?"* This is an important basic problem in robotic monitoring and data collection, and is one for which a variety of solutions have been proposed *e.g.*, [7, 5, 14].

Young-Ho Kim & Dylan A. Shell
Texas A&M University, College Station, TX, e-mail: (yhkim|dshell)@cse.tamu.edu
Colin Ho & Srikanth Saripalli
Arizona State University, Tempe, AZ e-mail: (colinho|srikanth.saripalli)@asu.edu

Linear least-squares estimation methods have been among the most successfully used in robotics for spatial interpolation and region sampling. By way of example, we include the recent work of Kemppainen *et al.* [7] in the Gaussian Process regression framework, and Elston *et al.* [5] via Kriging interpolation. Both employ equivalent minimum error-variance estimation techniques [11] that permit measured data to be interpolated in a way that takes into account a statistical description of spatial covariance [3]. Along with an estimate of values of interest at particular locations in the field, these methods also associate a measure of estimate uncertainty. Within the Kriging framework, the standard error measure is called the Kriging Variance (KV) or Kriging Error. It plays an important role for robotic adaptive sampling due to the fact that researchers (including, but not limited to [7] and [5]) have used it to select future sampling locations and to plan informative paths.

This paper assesses the role of pose uncertainty and measurement error on variance of such an estimated spatial field. We conducted a systematic empirical comparison of scalar fields reconstructed from measurements taken by robots using imperfect sensors.

1.1 Problem Statement

Two complementary aspects are investigated experimentally in this work:

Issue 1: Pose and measurement uncertainty — Classical formulations for spatial field estimation involve idealizations that may be ill-suited for robotic sampling. Specifically, standard formulations lack explicit consideration of measurement noise and position uncertainty, flying in the face of practical experience with real sensors. This paper includes a formulation which addresses these two aspects. Because significant prior robotic work employs formulations without modeling these forms of uncertainty explicitly (*cf.* [7, 5, 14]), we conducted an empirical evaluation of how estimates of the spatial field differ depending on whether the method employed considers these sources of uncertainty or not. Scale and severity of the uncertainty are important considerations too.

Issue 2: An alternative to KV — Yamamoto [12] argued that the traditional KV measure fails to measure local data dispersion appropriately because it is computed from a global description of spatial variance (the variogram) averaged over the whole estimate. To address this shortcoming he introduced a new measure called the Interpolation Variance (IV). It is computed as the weighted average of the squared differences between measured data and the interpolation estimate, which is intuitively analogous to the traditional expression for statistical variance. Several geostatistical papers have evaluated IV and compared it to KV [10, 13], but IV appears to be unknown in the robotics literature.

Figure 1 illustrates the difference in estimated field uncertainty when sample position uncertainty and measurement noise are factored into estimates of field uncertainty measured with KV. Without considering these aspects, field variance is underestimated. The figure on the right also shows the IV measure for comparison; IV offers a distinct and in some ways more informed estimate of the interpolation uncertainty.

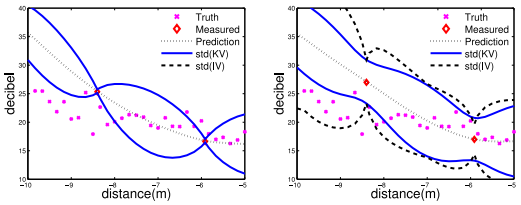

(a) Uncertainty ignored

(b) Pose & measurement error

Fig. 1 Richer treatment of site uncertainty is justified. The standard model (a) underestimates field variance compared to explicit treatment of uncertainty (b). Note also Interpolation Variance (IV) shown in (b).

The following are the paper's contributions:

- A unified Ordinary Kriging (OK) formulation with both pose uncertainty and measurement noise which explicitly separates sensing error from the variogram.
- A particle filter realization and implementation of the formulation.
- An empirical investigation of the impact of sample site uncertainty with respect to KV and IV. Most importantly, IV is an informative measure for selection of future measurement sites, which suggest a new set of adaptive sampling approaches.
- The demonstration of autonomous sampling site selection with single and multiple robots and a comparison KV and IV in terms of resulting performance.

Notation: Since no single definitive formalism has yet emerged for robotic sampling, we have elected to describe the results with terminology from the Kriging framework. The following motivated this choice:

1. Interpolation Variance has only been proposed within the Kriging framework.
2. Our work requires that we communicate our findings, including measured and estimated data, with geostatisticians.
3. Our applications involve robots estimating two or three dimensional fields. The family of variogram models have been informed by experience with the underlying physical processes (*cf.* covariance models in [9]).

Equivalency implies that the results hold for the Gaussian processes model as well.

1.2 Related Work

Measurement error is a common problem in robotics and it is widely recognized that most sensors are imperfect. In contrast, Cressie's classic geostatistics text [2] mentions that measurement error is usually implicitly included the nugget variance and that the practitioner often ignores the measurement error because it is typically considered smaller than the spatial variogram. When neither identified nor directly treated, micro-level variation of the spatial process is conflated with error introduced by imperfect sensing. The distinction is particular important when the variogram is intended to describe the intrinsic spatial variability of the statistical process, rather than its observation through a particular sensor. This is significant when multiple

sensors are available for a particular phenomenon and when sensor noise characteristics are estimated online.

Also, perfect robot pose information is not normally available. Chilés [1] was the first to consider position uncertainty in the Ordinary Kriging (OK) interpolator. He assumed a probability distribution over poses as the only form of error, naming this "attribute position error." Cressie and Kornak [3] take both position and measurement uncertainty into account in their Universal Kriging (UK) model. The UK variant is more general than the OK model, since it includes parameters to estimate strong trend in the underlying spatial process. The UK's generality comes at a price, however: interpolation error is not (numerically) simple to estimate. In fact, most existing robotics experiments employ OK [7, 14] over UK for this reason. Additionally, evidence presented in by Zimmerman *et al.* [15] shows that OK actually outperforms UK for realistic data since the latter is prone to over-fit. The OK formulation presented here is suited for robotics as it incorporates both pose uncertainty and measurement noise, is tractable when realized with a particle filter, and explicitly represents the sensor noise model.

In robotics, Kemppainen *et al.* [7] used Kriging without treatment of pose or measurement uncertainty in considering well-localized robots in an experimental environment. Zhu *et al.* [14] applied OK to find a optimal sampling strategy based on KV. Elston *et al.* [5] evaluated trajectory quality on the basis of a variogram model; the simulation did not consider measurement error.

As a measurement of interpolation error, KV does not depend on data values but only on the semivariogram model and data sample sites. Yamamoto [12] suggested that the Interpolation Variance (IV) would be well-founded because it measures the reliability of estimates by considering both the Kriging weights and data values, *i.e.,* it is a heteroscedastic measure. We are not aware of any evaluation of IV that considers measurement or position uncertainty.

2 Technical Approach

The purpose of Kriging is to estimate the values of a spatial random variable, Z, from sparse sample data. Ordinary Kriging, the most common type of Kriging in practice, assumes that the spatial statistical process that generates the random values can be characterized by an unknown mean [11]. An estimate $Z^*(x_0)$ is obtained from λ_i and samples $Z(x_i)$ via (1):

$$Z^*(x_0) = \sum_{i=1}^{n} \lambda_i Z(x_i), \quad \sum_{i=1}^{n} \lambda_i = 1. \tag{1}$$

The coefficients λ_i, called Kriging weights, show that the interpolated points are obtained from a distance weighted average of nearby measured points. The weights are computed in matrix form via a pseudo-inverse of (2):

$$\mathbf{A} = \begin{pmatrix} \gamma(x_1,x_1) & \cdots & \gamma(x_1,x_n) & 1 \\ \vdots & \ddots & \vdots & \vdots \\ \gamma(x_n,x_1) & \cdots & \gamma(x_n,x_n) & 1 \\ 1 & \cdots & 1 & 0 \end{pmatrix}, \tag{2}$$

$$\mathbf{b} = \begin{pmatrix} \gamma(x_1,x_0) \\ \vdots \\ \gamma(x_n,x_0) \\ 1 \end{pmatrix}, \quad \mathbf{A}\boldsymbol{\lambda} = \mathbf{b}. \tag{3}$$

The value depends on the semivariogram function γ which represents the strength of spatial relationships in the random field. It is defined as the square of the expected difference between values at different locations, *i.e.*, $\gamma(x,y) = [E(Z(x)) - E(Z(y))]^2$.

KV is given by $\sigma^2(x_0) = \sum_{i=1}^{n} \lambda_i \gamma(x_i,x_o) + \psi(x_o)$. Taken together, these yield the optimal OK prediction of unobserved values of the process when neither measurement noise nor position uncertainty are included.

2.1 Measurement error

Kriging Variance: The traditional Kriging approach usually treats the measurement noise as zero, *i.e.*, that the measured data represent true values. It is common for geostatistical texts (*e.g.,* [2, 11]) to add that if there is reason for a practitioner to suspect measurement error (described with variance, c_m) then the remedy is to subtract this value from the field Kriging variance, as so

$$\sigma_m^2(x_o) = \sum_{i=1}^{n} \lambda_m^{\,i} \gamma_m(x_i,x_o) + \psi(x_o) - c_m, \tag{4}$$

where γ_m is the variogram and ψ is a Lagrange multiplier. This is both counterintuitive and incorrect for imperfect measurement as typically considered in the robotic sampling context. Understanding why depends on two observations. Firstly, despite moving from idealized to noisy sensing, it has tacitly been assumed that the variogram is constructed using the same noisy sensor as employed for the sampling itself. Thus, the variogram no longer merely encodes the intrinsic spatial dependency of the underlying statistical process. The result is that γ_m is a variogram that includes measurement error and has been shifted by c_m automatically. Secondly, applying (4) removes the implicitly captured sensing variation in order to describe only the interpolation variance. It does not, therefore, yield a variance that estimates what would result if you were to place the sensor at the given location.

Instead, we define γ_m to be $\gamma_{truth} + c_m$, where γ_{truth} is calculated using a high-fidelity sensor or from a theoretical understanding of the spatial process. The resultant expression for the variance is

$$\sigma_\varepsilon^2(x_o) = \sum_{i=1}^{n} \lambda_{truth}^{i} \gamma_{truth}(x_i, x_o) + \psi(x_o) + c_m. \tag{5}$$

One is now no longer subtracting the measurement error from the prediction, which decreases the variance so it no longer matches measured data, but instead increasing the degree of uncertainty, as one might expect.

Interpolation Variance:

In contrast, interpolation variance estimates are affected by the variogram only indirectly. The IV computation is analogous to a form of variance where the probability density is replaced by OK weight λ. The choice of γ_m alters the computed λ, but the sensitivity to measurement noise is dominated by the local data variation. Clearly this latter effect has no impact on KV. Equation (6) shows IV in terms of λ and the variance of the input consistent with the other notation:

$$s_o^2 = \sum_{i=1}^{n} \lambda_i [z(x_i) - z^*(x_o)]^2. \tag{6}$$

2.2 Position uncertainty

To take position uncertainty into account within the Kriging framework, Chilés [1] proposed the following model. The measurement believed to be made at a point $\mathbf{x}_\alpha$ is actually performed at another position $\mathbf{x}_\alpha + \mathbf{u}_\alpha$, where $\mathbf{u}_\alpha$ is a random vector which contributes to the position error. Given probability density $p(\mathbf{u}_\alpha)$ and also the joint density $p(\mathbf{u}_\alpha, \mathbf{u}_\beta)$ for any pair of points, one may extend the OK equations (2)–(3) by randomizing vectors u_α over poses. This produces equations (7)–(8):

$$\begin{aligned} \tilde{\mathbf{x}}_\alpha &= \mathbf{x}_\alpha + \mathbf{u}_\alpha, \\ \tilde{\mathbf{A}}_{\alpha\beta} &= \iint \gamma_m(\tilde{x}_\alpha, \tilde{x}_\beta) p(u_\alpha, u_\beta) du_\alpha du_\beta. \end{aligned} \tag{7}$$

$$\tilde{\mathbf{b}}_{\mathbf{o}\alpha} = \int \gamma_m(\tilde{x}_\alpha, \tilde{x}_o) p(u_\alpha) du_\alpha. \tag{8}$$

The expressions for $\tilde{\mathbf{A}}_{\alpha\beta}$ and $\tilde{\mathbf{b}}_{\mathbf{o}\alpha}$ can be directly applied in (2)–(3) in place of $\mathbf{A}$ and $\mathbf{b}$. Employing this strategy propagates pose uncertainty (position variation) through to the field estimate and produces an increase in the interpolation variance. It is important to note that even if the distributions are unimodal and isotropic, the observed values may differ from the interpolator predictions at the associated maximum likelihood positions.

In practice, incorporating pose uncertainty in Kriging for real systems involves application of (7)–(8) in a representation which ensures the operations involved can be carried out tractably. We describe an efficient realization of Kriging via a particle filter. Additionally, this has the advantage of already being a practical representation of pose uncertainty, *e.g.*, see [6]. Particle filters approximate a probability density

through a set of weighted samples, $\omega_{(n)}$, drawn from the distribution being represented.

The following equations illustrate a Kriging computation, analogous to that described above, using the particle filter. Essentially, integrals have been replaced with summations, and appropriate weights used:

$$\tilde{\mathbf{A}}_{\alpha\beta} = \sum_{i=0}^{n} \sum_{j=0}^{m} \gamma_m(x_i, x_j)\, \omega_i\, \omega_j, \tag{9}$$

$$\tilde{\mathbf{b}}_{\mathbf{o}\alpha} = \sum_{i=0}^{n} \gamma_m(x_i, x_o)\, \omega_i, \tag{10}$$

where n and m correspond to the number of samples drawn from the distribution. Similar to the equations described above, once the relevant matrices are computed in this way, interpolation proceeds as before.

Our particular robots localize themselves using a particle filter based approach, so this treatment is particularly straightforward: Kriging inputs are computed via (9)–(10) on the localization particles and, once the relevant matrices are computed in this way, interpolation proceeds directly.

3 Experiments

Equipment & Experimental Scenarios

Two separate experimental scenarios were used. Both assess the estimated variance in scalar fields constructed from measurements taken by robots:

1. Sound and light, indoors: An iRobot Create robot, equipped with a Hokuyo URG-04LX-UG01 laser sensor and an Asus Eee PC 1005HA netbook, was given a map of our building; it moved around the environment using an adaptive particle filter to localize itself. The localized robot was given several target measurement positions along a linear corridor. A total of 44 measurement positions, each approximately 50*cm* apart, were provided as navigation goals. After the robot arrived at each goal position, it saved its current localization estimates by writing the particles representing the pose probability distribution to disk. The robot remained in place for 2 seconds, collecting data from three sensors: (1.) ambient sound volume in decibels via the onboard microphone on the netbook; (2.) light intensity measured in lux with a Phidgets "Precision Light Sensor" and (3.) ambient temperature via a Phidgets "Precision Temperature Sensor". We did not report values from the third sensor, as the building climate control made them uninteresting.

Figure 2 shows the second floor of our building. For simplicity the measurements we analyze further all come from a single corridor about 17*m* in length. Two sound sources emit the same continuous beeping sound at locations shown in the figure (the sound itself is a flute C7 note at 2000Hz.) The measurements were recorded at night so that the florescent lights in the ceiling were the sole light source.

The light and sound fields represent opposite extremes. Light intensity was essentially constant, except with a spatial period representing light spacing, the light sen-

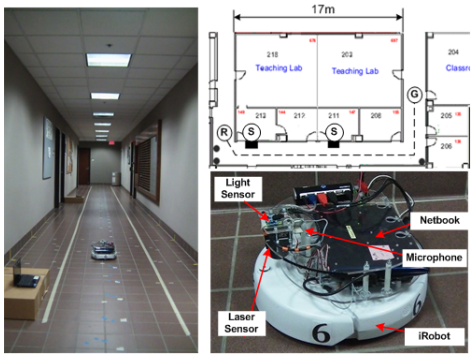

Fig. 2 Environment in which measurements were recorded: the robot navigates along the corridor from R to G, taking denote measurements. The two S symbols denote the positions of sound sources.

sor showed low variation, and the empirical variogram shows a short range spatial covariance. On the other hand, the measured sound volume had large fluctuations and a variogram illustrating longer range spatial structure.

The sensor noise variance c_m was estimated for each sensor, by computing the mean of variances computed from 200 values taken over 2 seconds at the same position. Variogram models were fitted passing through the origin from samples averaged over the 200 values in order to minimize sensor noise. We adopted the spherical variogram model (11), where n is the height of the jump of the semivariogram at the origin (known as the *nugget*), s^* is the limit of the variogram at infinite lag (the *sill*), and r is the distance at which this is first reached (*range*). Light within our environment is characterized as follows: $r = 2m$, $s^* = 580lux$, and the sensor has variance $c_m = 5$. Similarly sound has $r = 5m$, $s^* = 80$ decibels, and the sensor variance $c_m = 30$.

$$\gamma_{sph}(x_1, x_2) = s^* \cdot \left[\frac{3\|x_1 - x_2\|}{2r} - \frac{\|x_1 - x_2\|^3}{2r^3} \right] + n. \tag{11}$$

2. Interpolated H_2O quality indicators, from Lake Pleasant: An AUV equipped with a YSI 6600vs Sonde (see Fig. 3) took dense measurements within a $70m \times 100m$ window. These readings along with GPS positions were used as "ground truth" data on which we simulated sampling experiments as follows: readings at any point were generated from an OK interpolation with Gaussian noise added (σ^2 selected per YSI's listed measurement accuracy per quality indicator).

Samples, collected in the way described above, were then processed four different ways in order to consider the different treatments of sample site uncertainty:

1. The single best pose hypothesis with field measurement assumed to be error free (PBMB).
2. The single best pose hypothesis with field measurement treated as noisy (PBMU).

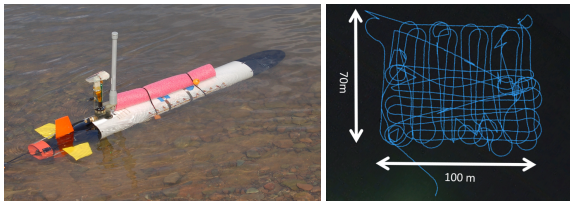

(a) Ocean Server Iver2 AUV. (b) Sampling pattern.

Fig. 3 Real world data were collected in Lake Pleasant, AZ. (−33° 51' 55.66"N, 112° 17' 45.01"W) by directing a AUV in (a) to trace the 3215m long "lawnmower" pattern in (b). Measurements of five different quality indicators were taken 3448 times at 0.5Hz over the 1hr 55min traversal.

3. The full pose distribution with measurement assumed to be error free (PUMB).
4. The full pose distribution with measurement treated as noisy (PUMU).

Instances involving PB used the single sample from the particle filter with maximum weight. The PU instances employed the whole distribution in order to pass the uncertainty onto the spatial field representation. The measurement treatment was analogous: variance of the sensor only being considered in MU cases.

Planners: Greedy KV- and IV-based Selection

In the data reported in Section 4, we provide the distance travelled as a robot attempts to reduce field uncertainty. In these cases, we considered a naïve planner in which the robot selects new sample locations by greedily picking the position in the field estimate with greatest uncertainty, sampling there, incoporating the data, and repeating the process. This uncertainty is measured using either KV or IV, as will be indicated.

4 Results

4.1 Comparison of KV and IV: Density Dependence

Figure 4 illustrates that KV depends on sample position but is independent of the observed data.[1] KV has the same shape between measurement sites no matter what the measured data actually are. However, IV considers different local data values and exhibits some degree of anisotropy in shape.

Figure 4 also shows that KV and IV have different characteristics and that these depend on the two different sensors. As the number of measured data increases, KV decreases smoothly. In contrast, IV can increase due to an implicit dependency on data variation. Thus, IV has the smaller value of the two estimates for sparse data, but with the increasing density, the tendency of both measures causes them to cross at a certain point. Figure 5 quantifies this effect for the sensors we considered: the light sensor results in an intersection for between 3 and 5 samples. For the sound sensor this occurs between 7 to 8 samples. Moreover, the data variance decreases

[1] Since the variogram is isotropic, the KV inherits this property too.

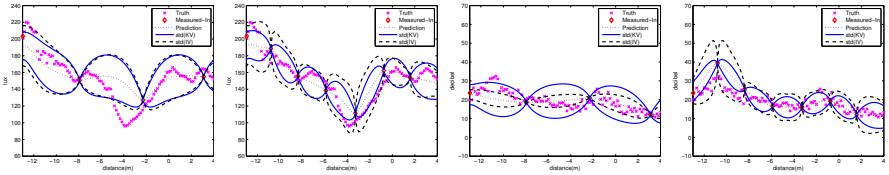

(a) Light, 4 samples (b) Light, 8 samples (c) Sound 4 samples (d) Sound 8 samples

Fig. 4 The two variance estimates have different behavior with regard to sample density.

as additional measurements are added. The evaluation of KV and IV for the two sensors in subsequent sections uses 4 inputs for light and 8 inputs for sound.

Fig. 5 Differing behavior as a function of data density. Analysis of variance (measured and estimated) as a function of sample density. (KV in red, IV in blue.)

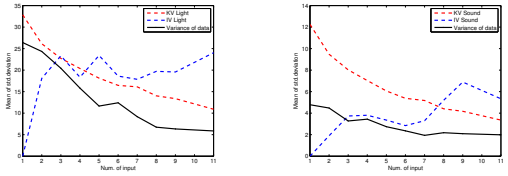

(a) Light sensor: Variances. (b) Sound sensor: Variances.

4.2 Measurement error

As the reported c_m values illustrate, the light sensor is considerably more sensitive than the sound sensor. To make clear the impact of measurement error on variance estimates, we show only the sound case.

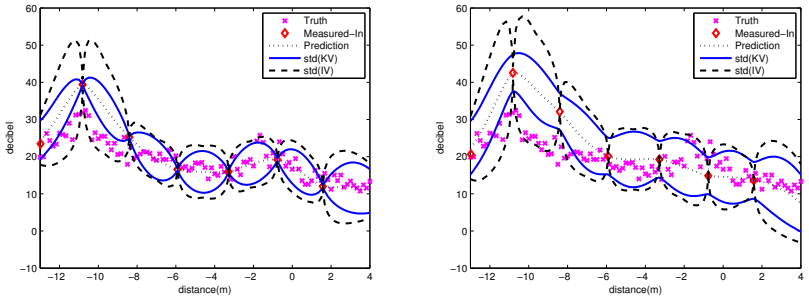

(c) PBMB Sound (d) PBMU Sound

Fig. 6 Comparison of the measurement error at 8 points.

Figure 6 illustrates several properties: IV equals zero at the measured points, while KV is increased by c_m at the measured points. Moving beyond the measured points, both increase approximately linearly with increasing noise. Figure 6 shows clearly how IV underestimates at measured points. The rate that IV increases reflects unexpected uncertain data unlike the KV which merely considers distance between measurements.

Although IV vanishes at measurement points, it compensates when multiple sensed values are close to one another since the sensor noise inherently limits the quality of prediction from the local neighborhood. The measurement error is implicitly represented by IV, which relies on the increasing numbers of measurements, to truly reflect the data spread. We suggest that IV will capture c_m when multiple measurements are taken from the same location. However, IV must be taken as the *envelope* computed from all measurements and one should not compute IV from the mean.

4.3 Position uncertainty

Figure 7 shows the effect of position uncertainty for each of the sensed fields. Neither, KV nor IV are zero at the measured points but both increase linearly with IV increasing more than KV. Severe position uncertainty degrades interpolation because $\tilde{A}_{\alpha\beta}$ represents a convolution weight. It is worth observing that KV increases most near measurements. On the other hand, IV increases variance everywhere else. Unlike the previous case, the interpolated Kriging may no longer be a punctual interpolator, *i.e.,* the estimate may not coincide with the measurement.

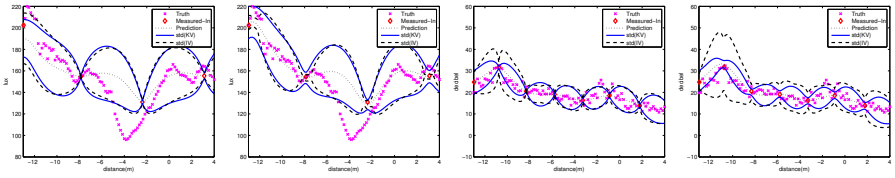

(a) PBMB Light. (b) PUMB Light. (c) PBMB Sound. (d) PUMB Sound.

Fig. 7 Position uncertainty: light with 4 inputs, sound with 8 inputs.

4.4 Effects of both forms of sample uncertainty

It is worth considering both measurement error and position uncertainty together, as it best reflects the reality of most robotic sampling scenarios. Figure 8(a) shows the PBMB case (no uncertainty modelled) for the light sensor. The KV, however, underestimates the variance in some locations, partly because the light sensor has a small variance. Figure 8(c) shows the underestimation of variance (observable by comparison to the ground-truth data). By considering uncertainty, Figure 8(b) shows that KV does not fully resolve the underestimation problem, while IV appears to be

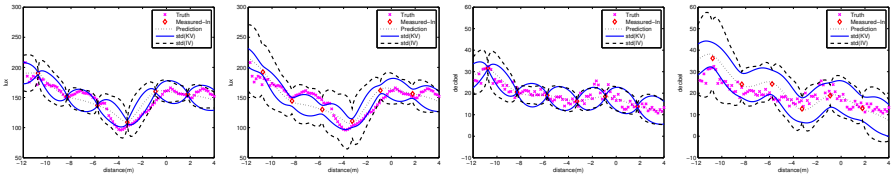

(a) PBMB Light. (b) PUMU Light. (c) PBMB Sound. (d) PUMU Sound.

Fig. 8 Both variance measures for 8 light and sound readings.

a better estimate. However, it may overestimate in cases with small scale sensor variance like the light sensor. In Figure 8(d), the KV does not show the variation of sensor values. However, IV shows good estimates for the noisy sound sensor, and the PUMU IV case shows the uncertainty near the measurements.

In many cases, the absolute variance estimate is used to select navigation goals. While both forms of variance can be used, it is important to note that the sensor variance can influence the appropriate choice of goal.

4.5 Planner effectiveness of for adaptive sampling

Greedy planners based on KV and IV can both be effectively used for autonomous sampling since either decreases (both) measures of variance. Fig. 9 summarizes the experiments conducted with a single robot indoors for estimating the light intensity field. Selecting targets with KV results in an approximately even spacing between samples, as shown in inter-sample distance histogram inset in Fig. 9(a). In contrast, the insets in Fig. 9(b) show how the heteroscedastic property causes different spacing. The right inset of Fig. 9(b) also illustrates the importance of employing the envelope scheme when there are multiple measurements at a particular site or even very nearby sample sites.

Figure 10 shows data from multi-robot tests with KV or IV for sample site planning. The plots show that the two robots are able to sample with efficiency that is almost twice that of the individual robots, if we consider the rate of decrease in total field variance.

4.6 Data-dependence on planner performance

While KV is widely known and a frequently used measure, IV has complementary aspects as a candidate for robotic sampling applications. Fig. 11 shows that this is observation is true in both scenarios. For light indoors, the robot covers a distance $\pm 25\%$ shorter to collect the same number of samples as a KV planning robot. In a given time interval, the IV method collects $\pm 45\%$ more samples than KV. The table below shows Lake Pleasant data.

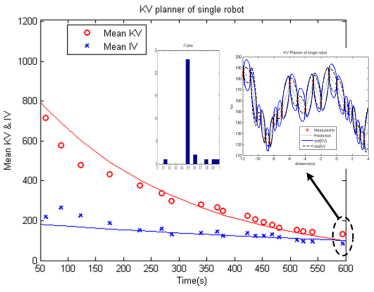

(a) KV-based planner (first 600 seconds).

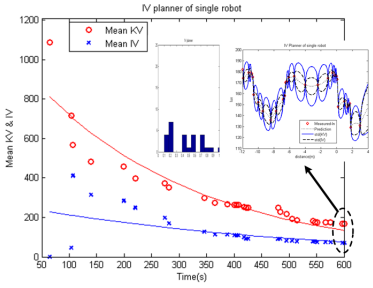

(b) IV-based planner (first 600 seconds).

Fig. 9 Autonomous sampling drives variance measures down. The sample position selected as either (a) maximal KV, or (b) maximum IV position. The vertical axis represents the mean KV and IV, the horizontal axis denotes time. The inset histogram shows the distribution of inter-sample distances. The right inset is the resultant light field (in lux) along with variance estimates.

	IV planner	**KV planner**
B/G Algae	$(3467.7, 37104)$	$(3186.8, 26543)$
Chlorophyll	$(0.4117, 0.0015)$	$(0.4321, 0.0013)$
pH	$(0.1509, 0.0008)$	$(0.1687, 0.0004)$
Temperature	$(0.1939, 0.0001)$	$(0.2071, 0.0002)$

(Left) Reported statistics are (μ, σ) of MSE between estimated field and ground-truth values, for $n = 10$ independent trials.

Table 1 Application of IV and KV planner in Lake Pleasant, AZ. The results compare 10 cross-validation trials, showing the mean and standard deviation values. The estimated field uses 30 samples at point A in Fig. 11(b).

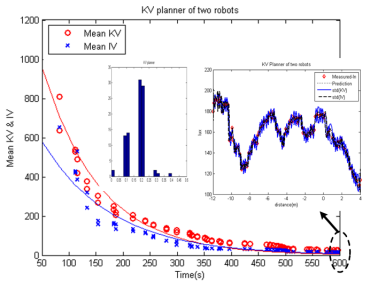

(a) KV for selecting sample sites.

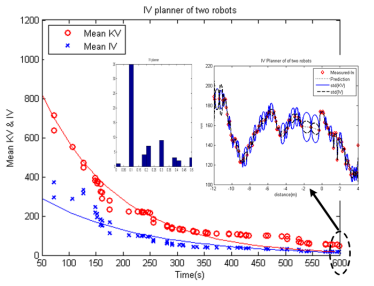

(b) IV for selecting sample sites.

Fig. 10 A multi-robot demonstration of KV and IV for sampling. Two robots sample light and temperature concurrently.

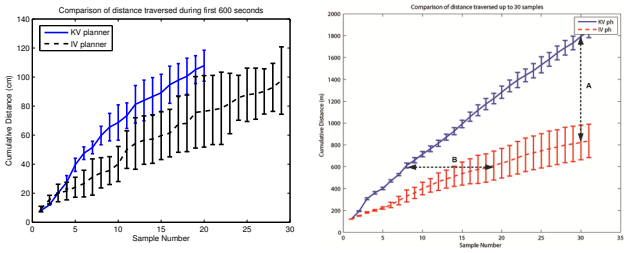

(a) Indoor Scenario (Light) (b) Lake Scenario (pH)

Fig. 11 Minimizing IV collects more data for a given duration than KV, and travels shorter distances to collect the same number of samples. Data are means and variances from ten separate trials of each greedy planner.

(Right) Cross-validation analogous to Table 1, but with an upper bound on distance traversed rather than a fixed number of samples. The s value denotes number of samples used.

	IV planner	**KV planner**
B/G Algae	$(3467.7, 37104)$ $s = 30$	$(3236.7, 29308)$ $s = 9$
Chlorophyll	$(0.4084, 0.0014)$ $s = 25$	$(0.5212, 0.0048)$ $s = 8$
pH	$(0.1786, 0.0024)$ $s = 18$	$(0.1796, 0.0008)$ $s = 8$
Temperature	$(0.1939, 0.0001)$ $s = 30$	$(0.2007, 0.0003)$ $s = 5$

Table 2 Application of IV and KV planner in Lake Pleasant, AZ. The results are mean and standard deviation of MSE measured cross-validation between estimated field and ground-truth values for 10 trials. The estimated field are for samples drawn within the first 600m in B of Fig. 11(b).

5 Conclusion

In robotics, Kriging spatial interpolation is generally used without considering measurement error and position uncertainty. However, our data demonstrate that doing so may result in underestimation of interpolation error. This paper shows how both of these forms of uncertainty can be easily incorporated into the standard OK interpolator, allowing the uncertainty to reflect itself as an increase in variance of the underlying field. Our physical robot data show that measurement error and position uncertainty do affect the error estimates for both KV and IV. Our empirical assessment resulted in several observations about the models variance and their relationship. For example, data density affects both KV and IV, but in opposite ways; we first postulated and then determined that a cross-over point exists. Secondly, even though IV equals zero at measured points, employing the envelope of the variance function is natural near multiply measured points. Finally, sample site uncertainty not only changes the variance, but can also result in a non-punctual interpolator.

While KV is widely known and frequently used traditional measure, IV is a new candidate for robotic sampling applications which has several aspects that are complementary to KV. We have shown the feasibility of IV as a measure of uncertainty by incorporating it in a demonstration of an autonomous robot system which adaptively samples light and temperature in our building. When compared to KV, selecting points to greedily minimize IV collects more data for a given duration and travels shorter distances to collect the same number of samples. Additionally, IV appears no less suited to multi-robot applications than KV in terms of potential sampling speed-up.

References

1. Chilés, J.P.: How to adapt Kriging to non-classical problems: three case studies. In: M. Guarascio, M. David, C. Huijbregts (eds.) Advanced Geostatistics in the Mining Industry, pp. 69–90. D. Reidel Publishing Co., Holland (1993)
2. Cressie, N.: Statistics for Spatial Data. John Wiley and Sons, Ltd, New York, NY, USA (1993)
3. Cressie, N., Kornak, J.: Spatial Statistics in the Presence of Location Error with an Application to Remote Sensing of the Environment. Statistical Science pp. 18(4), 436–456 (2003)
4. Dolan, J.M., Podnar, G.W., Stancliff, S., Low, K.H., Elfes, A., Higinbotham, J., Hosler, J., Moisan, T., Moisan, J.: Cooperative Aquatic Sensing Using the Telesupervised Adaptive Ocean Sensor Fleet. In: Proceedings of the SPIE Conference on Remote Sensing of the Ocean, Sea Ice, and Large Water Regions. Berlin, Germany (2009)
5. Elston, J., , Stachura, M., Frew, E., Herzfeld, U.: Toward Model Free Atmospheric Sensing by Aerial Robot Networks in Strong Wind Fields. In: Proceedings of the International Conference on Robotics and Automation (ICRA09). Kobe, Japan (2009)
6. Fox, D.: KLD-sampling: Adaptive particle filters. In: T.G. Dietterich, S. Becker, Z. Ghahramani (eds.) Advances in Neural Information Processing Systems (NIPS-14), pp. 713–720. MIT Press, Cambridge, MA, USA (2001)
7. Kemppainen, A., Makela, T., Haverinen, J., Roning, J.: An Experimental Environment for Optimal Spatial Sampling in a Multi-Robot System. In: Proceedings of the International Conference on Intelligent Autonomous Systems. Baden-Barden, Germany (2008)
8. Kumar, V., Rus, D., Sukhatme, G.S.: Networked Robots. In: B. Siciliano, O. Khatib (eds.) Springer Handbook of Robotics, chap. 41. Springer-Verlag Heidelberg (2008)
9. Rasmussen, C.E., Williams, C.K.I.: Gaussian Processes for Machine Learning. MIT Press, Cambridge, MA, USA (2006)
10. da Rocha, M.M., Yamamoto, J.K.: Comparison Between Kriging Variance and Interpolation Variance as Uncertainty Measurements in the Capanema Iron Mine, State of Minas Gerais-Brazil. Natural Resources Research pp. 9(3), 223–235 (2000)
11. Webster, R., Oliver, M.: Geostatistics for Environmental Scientists. John Wiley and Sons, Ltd, New York, NY, USA (2007)
12. Yamamoto, J.: An Alternative Measure of the Reliability of Ordinary Kriging Estimates. Mathematical Geology **32**(4), 489–509 (2000)
13. Yamamoto, J.K., da Rocha, M.M.: Properties and Applications of the Interpolation Variance Associated with Ordinary Kriging Estimates. In: Proceedings of the International Symposium on Spatial Accuracy Assessment in Natural Resources and Environmental Sciences, pp. 70–77. Shanghai, China (2008)
14. Zhu, X., Yu, J., Ren, S., Wang, X.: Near-optimal Collecting Data Strategy Based on Orinary Kriging Variance. In: Proceedings of OCEANS, pp. 1–6. Sydney, NSW, Australia (2010)
15. Zimmerman, D., Pavlik, C., Ruggles, A., Armstrong, M.P.: An Experimental Comparison of Ordinary and Universal Kriging and Inverse Distance Weighting. Mathmatical Geology **31**(4), 375–390 (1999)