

Adaptive Multiscale Sampling in Robotic Sensor Networks

Vadiraj Hombal, Arthur Sanderson and D. Richard Blidberg

Abstract—This work focuses on the observation of environmental phenomena that occur as spatial distributions in two and three dimensions, using sensor-enabled mobile vehicles (ground, air or undersea). Algorithms to guide an adaptive exploration of a given region through systematic choice of sampling locations under the constraints imposed by vehicles are presented. Variation sensitive multiresolution sample distributions are achieved through an iterative variation sensitive estimation of the unknown process.

I. INTRODUCTION

Mobile robots are increasingly employed in the observation and monitoring of complex processes on the ground, in the air, and undersea [3], [7], [8], [11], [13]. An example is the observation and mapping of harmful algal blooms in order to predict occurrence and trends that impact biological processes and human health [2].

Advances in sensor technologies and Autonomous Underwater Vehicles (AUVs), have been incorporated in distributed sensor networks enabling pervasive *in-situ* observation of such processes in a wide range of spatial and temporal sampling resolutions [8]. Effective deployment such resources requires integration of the sampling objectives with constraints such as the fixed number of vehicles available, and the limited energy available for sensing.

The objective of environmental observation is to extract the synoptic structures (space-time snapshots) of the phenomena of interest in order to make effective predictive and analytical characterizations about it [12]. Section II presents a formulation the objective and discusses the issues involved in its realization.

Introduction of sensor nodes on AUVs enables adaptive choice of sensing tasks and dynamic redeployment of monitoring resources in response to new observations and predictions. This requires a field model to integrate intermediate measurements into an estimate of the underlying function and guide the selection of subsequent measurements. For complex oceanographic processes, typically there are no known models. Even when such models are available, the large computational costs associated with them, prohibit their use for online decision making on a mobile platform. Instead, surrogate models must be employed to drive adaptivity. A principle contribution of this paper is the introduction of

multiscale surrogate models for use in such applications. These are based on a model that implements a hierarchical decomposition of the error space (Section III). Following this, the MSAS algorithm, which implements the corresponding sequential decomposition and achieves variation sensitive sampling distribution through variation sensitive modeling, is presented in Section IV.

The use of mobile sensor platforms for observation of spatio-temporal environmental processes involves a decision making process that trades-off the amount of information gathered with vehicle and process constraints. These are presented in Section V.

II. RECOVERY OF SYNOPTIC FEATURES

The principle objective of this work is the recovery of synoptic features of an environmental process $f(x, t)$, $x \in \Omega$, $t \in [0, T]$ from its space-time point samples in a domain Ω of area $A \triangleq |\Omega|$, over T days. The goal is to obtain samples at a spatial resolution that allows extraction of a series of spatial snapshots over Ω at a rate appropriate to the variation of the process in time. Towards this, nominal values for spatial resolution, $\Delta x > 0$, and temporal resolution, $\Delta t > 0$, are assumed.

In each interval Δt the goal is to collect the *best* possible data $\{f(X_s), X_s; X_s = \{x_i\}_{i=1}^k, x_i \in \Omega\}$ given the dynamic and energy constraints of the vehicles. The vehicles have a maximum speed of V , are subject to non-holonomic constraints and have a limited energy E . Within each interval Δt , the quality of the data can be measured by:

$$I(X_s) = \sum_{x \in \Omega} \left| f(x) - \hat{f}(x; X_s, f(X_s)) \right|^2 \quad (1)$$

where the dependence of measurements on time t , has been suppressed for ease of notation.

A principal factor in the optimization of this objective is the choice of the approximating model \hat{f} . If the model is well validated and represents the primary variable of interest, then sample distribution may be designed in the classical design of experiments framework [5]. Depending on the chosen metric, such sample designs are distributed so as to exercise the variability in the model parameters, their predictions [4], or perhaps the models themselves [9]. If the models are linear in their parameters then the sample design is independent of the measurements $f(X_s)$, and the design can be selected a-priori so as to exercise the variability in the respective metric. In such cases, the observed synoptic feature itself derives entirely from the model rather than the sample distribution.

However, typical sampling missions are exploratory in nature and are planned to develop an observation base of a

Funding was provided by grant number IIS-0329837 from the National Science Foundation and by the Office of Naval Research (ONR) under grants #N00014-04-1-0264, #N00014-03-C-0109, #N00014-02-1-008.

Vadiraj Hombal, Dept. of Electrical, Computer & Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY hombav@rpi.edu

Arthur Sanderson, Dept. of Electrical, Computer & Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY. sandea@rpi.edu

D. Richard Blidberg, Autonomous Undersea Research Institute (AUSI), Lee, New Hampshire blidberg@ausi.org

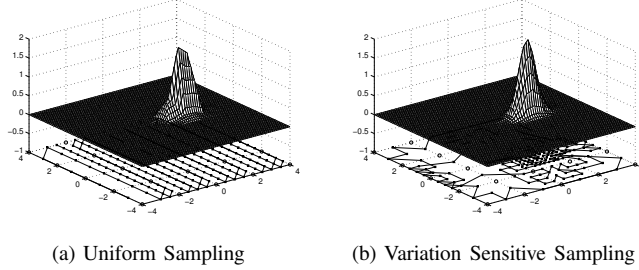


Fig. 1: Coverage and Resolution of Sample Distributions

process (for e.g., biological) for which no prior information exists. In such cases, a sample design is expected to be the observational basis that provides empirical evidence for features that exist in the domain. Towards this, sample distributions must facilitate extraction of synoptic features both through the coverage of the domain and resolution of the features at appropriate scales. The coverage of a sample distribution X_s may be defined as the area of its convex hull $A = |X_s|$. As $A \rightarrow |\Omega|$, errors due to extrapolation are minimized.

In such problems, the standard design for AUV platforms is uniform grid sampling (Fig 1a). Willcox [14] has presented guidelines for the design of uniform sample distributions under process and resource constraints. The key design variables in such cases are the length of the transects and the separation between them. The typical condition of such sample designs, is that while regions void of features are over sampled, feature regions are under sampled. Allen et al., [1] cite that the coarse spatial resolutions of uniform sampling rivals the size of the features rather the variations in it. Further, coverage can only be achieved at the cost of resolution in such designs. Literature in oceanography reports 40 – 50% errors due to extrapolation [12] that results from lack of coverage. In addition, if the measurement costs per sample are not negligible, as is typically the case in oceanographic sampling, then uniform sampling represents an inefficient sampling regime.

In order to achieve a balance between these conflicting requirements, this work proposes the use of variation sensitive sampling (Fig 1b) that attempt to balance the above mentioned requirements. Since, the underlying function is unknown, mechanisms which generate such sample designs, must incorporate incremental knowledge available through measurements. The principle focus of this work is an iterative solution for the estimation of the unknown process that is needed for this purpose.

III. HIERARCHICAL RADIAL BASIS FUNCTIONS

Hierarchical Radial Basis Functions (HRBF) implement a multilayered hierarchy of Radial Basis Function Networks (RBFN) of different scales to generate a coarse to fine representation of a given function through a hierarchical multiscale decomposition of the approximation error space [6].

Such decompositions are then employed towards the sequential coarse to fine modeling of an *unknown* function and adaptive coarse to fine sample distribution.

A. Hierarchical Decomposition

An M -level hierarchical decomposition of a function $f : \Omega \rightarrow \mathbb{R}$, $\Omega \in \mathbb{R}^D$ may be represented in terms of approximations: $e_k = \hat{e}_k + e_{k+1}$, $k = \{1 \dots M\}$, where $e_1 = f$ and \hat{e}_k is the k^{th} layer approximation to the error, e_k , in the previous layer, and $e_{k+1} = e_k - \hat{e}_k$ is the corresponding approximation error of the current layer. Alternately, the decomposition may be written as:

$$f = \sum_{k=1}^M \hat{e}_k + e_{M+1} \triangleq \hat{f}_M + e_{M+1}$$

where \hat{f}_M is defined as the M -level approximation to function f , and $e_{M+1} = f - \hat{f}_M$ is the error due to the M -level approximation, \hat{f}_M . In general, e_k is the approximation error due to the $(k-1)$ level hierarchical approximation of f : $e_k = f - \hat{f}_{k-1} = f - \sum_{j=1}^{k-1} \hat{e}_j$. Thus each layer, k , approximates the error space due to the $(k-1)$ level approximation of the function, f at scale k .

In HRBF, at the k^{th} level of hierarchy, the error at point $\mathbf{x} \in \Omega$, $e_k(\mathbf{x})$, is approximated by an RBFN of N_k basis functions $\phi_{k,j}(\mathbf{x}) = \phi(\|\mathbf{x} - \mathbf{c}_{k,j}\|; \sigma_k)$ of scale $\sigma_k > 0$, which are centered at locations $C_k = \{\mathbf{c}_{k,j} \in \Omega, j = 1, 2, \dots, N_k\}$:

$$\hat{e}_k(\mathbf{x}) = \sum_{j=1}^{N_k} w_{k,j} \phi_{k,j}(\mathbf{x}) = \Phi_k(\mathbf{x}) \mathbf{w}_k$$

where, $\Phi_k(\mathbf{x})$ is the *interpolation* matrix whose N_k columns span the subspace represented by the k^{th} layer and \mathbf{w}_k is the corresponding vector ($N_k \times 1$) of approximation parameters. The M -level multiscale approximation may be written as:

$$\hat{f}_M(\mathbf{x}) = \sum_{k=1}^M \Phi_k(\mathbf{x}) \mathbf{w}_k = \Phi(\mathbf{x}) \mathbf{w} \quad (2)$$

Equation (2) represents a multi-scale approximation model that allows for *synthesis* of f in terms of a chosen basis set. In regularization theory such multiscale hierarchical models derive from the choice of stabilizer, which also yields the form of the radial bases [10]. In this work, Gaussian bases, which have a quasi-local support are assumed.

With the form of the radial bases fixed, the quality of approximation in (2) depends on the choice of the approximation parameters, and the structural parameters. In this work, the structural parameters are set so as to form a pre-defined partition of the input domain and the approximation parameters are set so as to optimize a chosen approximation metric. This setting of the structural parameters according to the geometry of the problem domain and independent of the approximation parameters yields a scale-space characterization of the function.

1) HRBF Structural Parameters and the Analysis Grid:

In this work the structural parameters of HRBF are set according to an ordering imposed by a hierarchical *analysis grid* defined on the problem domain, Ω . Each layer of the analysis grid corresponds to a layer in the HRBF. The number and position of nodes at the corresponding layer in the grid determine the number of basis functions, and the locations of the centers in each HRBF layer. Further, the scale parameter is set according to the density of the nodes in that layer.

Each layer of the analysis grid is a dyadic partition of the previous layer and the nodes correspond to the intersections of the. An example of an analysis grid defined over a domain $\Omega = [a, b] \in \mathfrak{R}$ is shown in Figure 2. The construction of the grid begins with the setting of N_1 , the number of nodes in layer 1, including two on the domain boundary. Nodes in subsequent layers are generated according to a *refinement* procedure that allows for a k^{th} level node $c_{k,j}$ to be expanded in terms of its child nodes, $C_{k+1}^{(k,j)}$ in level $k+1$:

$$C_{k+1}^{(k,j)} = \begin{cases} \{c_{k+1,2j-1}, c_{k+1,2j}\} & j = 1 \\ \{c_{k+1,2j-2}, c_{k+1,2j-1}, c_{k+1,2j}\} & 1 < j < 2^k + 1 \\ \{c_{k+1,2j-2}, c_{k+1,2j-1}\} & j = 2^k + 1 \end{cases} \quad (3)$$

where, location of the left child, $c_{k+1,2j-2} = \frac{1}{2}(c_{k,j} + c_{k,j-1})$, middle child, $c_{k+1,2j-1} = c_{k,j}$ and right child, $c_{k+1,2j} = \frac{1}{2}(c_{k,j+1} + c_{k,j})$. Also, $\rho_{k+1} = \frac{1}{2}\rho_k$. Since each higher layer is obtained through partition of the previous layer, $C_1 \subseteq C_2 \subseteq \dots \subseteq C_{M-1} \subseteq C_M$. As is discussed in subsequent sections, such a grid allows for inter-layer continuity in scale that is useful in sequential selection of critical bases of reconstruction.

Each level k , consists of $N_k = 2^k + 1$ centers, $C_k = \{c_{k,\cdot}\}$, including two on the boundary of the domain. Further, each node in the k^{th} layer is separated from its neighbour by $\rho_k = \rho_1/2^{k-1}$ and located at $c_{k,j} = a + \rho_k(j-1)$, $j = 1, 2, \dots, 2^k + 1$. The $c_{k,\cdot}$ bases in HRBF are centered at the location of the $c_{k,\cdot}$ node on the grid, with a scale: $\sigma_k = \alpha * \rho_k$. $\alpha > 0$, is a parameter that provides for a degree of freedom in the adjustment of scale. Thus, σ_k progressively decreases as k increases. This ordering on σ_k facilitates a coarse-to-fine decomposition of the function such that coarse scale trends in f are modeled in the initial layers and fine scale features are modeled in successive layers of the model. Extensions to higher dimensions follow directly. If the domain of interest is $[a, b]^D$, then starting with $N_1 = 3$ in each dimension, each layer of the analysis grid contains $N_k = (2^k + 1)^D$ centers, $C_k = \{c_{k,\cdot}\}$. Further, each node in the k^{th} layer is separated from its neighbour by $\rho_k = (b-a)/2^k$.

IV. MULTI SCALE ADAPTIVE SAMPLING

The Multi-Scale Adaptive Sampling (MSAS) algorithms presented in this section traverse the scale-space grid and arrive at a sparse representation of an a-priori *unknown* function through sequential discovery of significant bases in the grid.




Fig. 2: Analysis Grid

The algorithm starts with an initial coarse sample distribution and an initial model consisting of bases at a coarse scale. It then implements an iterative feedback procedure to systematically select nodes on the grid for subsequent sample refinement and model update. By linking sampling resolution and localization to basis scale and location, MSAS achieves sequential variation sensitive sample distribution via sequential variation sensitive modeling.

In each iteration, the selection of a node for refinement and modeling is based on criteria that use information contained in residues, which represent a sampling of the true approximation error in the current function estimate. Hence, the MSAS algorithm is an adaptive procedure for the refinement of existing sample distribution based on a-posteriori estimates of an adaptivity measure.

A. Mechanics of the MSAS algorithm

A procedural outline of the algorithm is presented below:

1. **select node** from the set of allowable candidate nodes
2. **sample function** at child nodes of the selected node
3. **add the basis** of the selected node to the model
4. **add children** of the modeled node to the set of allowable candidate nodes
5. **repeat** till some termination criteria is met

Procedurally, steps 1, 2 and 3 of the MSAS outline implement a *select-refine-model* cycle. The mechanics of the sequential MSAS algorithm which controls its trajectory through the scale space grid is governed by two key ideas:

1. In each iteration, the set of **allowable candidate nodes**, C_A , consists of all the un-modeled children of previously modeled nodes. This set represents the bases, ϕ_A , that are allowed to enter the model in the next iteration. From this set of allowable candidate nodes a subset of nodes is selected, according to the selection criteria described in section IV-D, for subsequent modeling.

2. A selected node is refined by sampling at its child node locations *before* it is modeled.

These ideas are illustrated in Figure 3. In Fig 3a, nodes C_1 correspond to previously modeled nodes, ϕ_1 which correspond to $\hat{f}_1 = \Phi_1 \mathbf{w}_1$. Nodes C_2 , which form the set of all un-modeled children of the previously modeled nodes, form the set of allowable candidate nodes for the next iteration. Suppose node $c_{2,2}$ is selected for modeling in the next iteration. The selected node $c_{2,2}$ is first refined by taking measurements at the locations of the child nodes $\{c_{3,2}, c_{3,3}, c_{3,4}\}$,




Fig. 3: Mechanics of the MSAS Algorithm

and is then introduced into the model: $\hat{f}_2 = \hat{f}_1 + \Phi_{2,2}w_{2,2}$. Modeling of the selected node $c_{2,2}$ introduces the child nodes into the set of allowable nodes for the next iteration. Thus, as seen in Fig 3b, the set of allowable candidate nodes is $\{C_1 \setminus c_{2,2}\} \cup \{c_{3,2}, c_{3,3}, c_{3,4}\}$.

From the modeling perspective, the above mechanics has two important consequences:

1. Allowing only un-modeled children of previously modeled nodes to enter the model ensures that each basis enters the model only once. This ensures that the iterative algorithm does not get stuck in a recursive loop in searching through the scale-space.
2. Since new bases can only enter the pool of allowable candidate nodes through refinement of parent nodes, no child node is added to the model before at least one of its parent nodes is modeled. This imposes inter-generational continuity (in both scale and space) on the bases and thus ensures that the model descends scale space grid in order.

The location of the un-modeled children of previously modeled nodes corresponds to all the previously sampled locations. Thus, in each iteration, all previously sampled locations are eligible for refinement. However the resolution of sample refinement is limited by the corresponding the scale of allowable bases. As a consequence of inter-generational continuity of bases, a region in the domain is sampled at a given resolution ρ_k , only after its neighbourhood has been sampled at resolution ρ_{k-1} .

As a consequence of the grid based regime, if a child node of a selected node is already modeled or sampled, it is not sampled again. Also, if a node is selected after both its neighboring nodes $c_{.,j-1}$ and $c_{.,j+1}$ have been selected, then no new samples need be taken. In such an iteration only a basis need be introduced into the model and only the basis corresponding to the middle child is newly added to the set of allowable nodes, since the other child nodes would have already been introduced by the neighboring nodes.

B. Maximum Sampling Resolution and Sample Budget

The mechanism described above allows for traversal of the scale space grid down to an arbitrary grid depth with arbitrarily high sampling resolution via refinement. However in practice, the maximum number of samples allowed and the maximum sampling resolution are limited due to cost of sampling and limitations of the sampling platform. Also, in

environmental applications, high sensor costs typically limit the available sample budget and hence the maximum number of samples allowed is fixed prior to mission design.

Without loss of generality, the maximum allowable sampling rate is assumed to correspond to the highest resolution of the grid and the maximum sample size, N_{\max} , is assumed to be lesser than 2^M , where M is the highest grid level. Also, the M^{th} level nodes are modeled *without* further refinement and therefore do not spawn new child nodes.

C. Modeling

In MSAS, since the refinement of a selected node *precedes* its modeling, the number of samples (N) is always more than the number of modeled nodes (n) and the resulting linear system represents an approximation rather than interpolation:

$$\Psi_i^*(X_i)\mathbf{w}_i^* = f(X_i) - \hat{f}_{i-1}(X_i) = e_i(X_i)$$

where, $X_i = X_{i-1} \cup C_R$ is the set of all the $N = |X_i|$ sampled points including nodes C_R obtained from refinement of the selected nodes C_i^* . $n = |\phi_*|$ is the number of new bases ϕ_i^* selected to be added to the model in the i^{th} iteration. In order to bootstrap the algorithm, the MSAS is initialized by selecting all layer 1 nodes, C_1 . Thus $C_1^* = C_1$. Refinement of these nodes in samples $f(X_1)$, at the location of all the layer 2 nodes, C_2 . Thus $X_1 = C_2$. The MSAS algorithm is presented below:

-
- 1) $i \rightarrow 1$
 - 2) set N_{\max}
 - 3) initialize: $e_i \rightarrow f$, $\hat{f}_{i-1} \rightarrow 0$, $X_{i-1} \rightarrow \emptyset$
 - 4) bootstrap: $C_A \rightarrow C_1$, $C_i^* \rightarrow C_1$
 - 5) $C_R \rightarrow \text{REFINE}(C_i^*)$
 - 6) $X_i \rightarrow X_{i-1} \cup C_R$
 - 7) analysis: $\Psi_*(X_i)\mathbf{w}_i^* = e_i(X_i)$
 - 8) synthesis: $\hat{e}_i \rightarrow \Phi_i^*\mathbf{w}_i^*$, $\hat{f}_i \rightarrow \hat{f}_{i-1} + \hat{e}_i$, $e_{i+1} \rightarrow f - \hat{f}_i$
 - 9) if $|X_i| \geq N_{\max}$ terminate.
 - 10) $C_A \rightarrow (C_A \setminus C_i^*) \cup C_R$
 - 11) $C_{i+1}^* \rightarrow \text{SELECT}(C_A, e_{i+1}(C_A))$
 - 12) $i \rightarrow i + 1$
 - 13) goto 5
-

D. Selection Criteria

As a consequence of the bootstrapping procedure described in the previous section, in each iteration, information about the error in the i^{th} - level approximation \hat{f}_i is available as residues $e_{i+1}(X_i) = f(X_i) - \hat{f}_i(X_i)$ at all the previously sampled locations, X_i . These correspond to the location of all nodes in the set of allowable candidate nodes A_C . In MSAS algorithm, the information contained in these residues is directly employed in each iteration to select a node, c_{i+1}^* from the set A_C such that, the resulting error in approximation is systematically reduced. Two selection criteria are presented: 1) greatest residue and 2) greatest error correlation:




Fig. 4: Scale Space Tree




Fig. 5: Samples: Top MSAS-GR, Bottom MSAS-GC

1) *MSAS-GR*: In MSAS algorithm based on greatest residue (MSAS-GR), the node corresponding to the greatest residue is selected from the set of allowable candidate nodes for modeling in the next iteration:

$$\text{selected node, } c_{i+1}^* = \arg \max_{c_\lambda \in A_C} (|e_{i+1}(c_\lambda)|)$$

2) *MSAS-GC*: Let $\tilde{e}(c_\lambda, X_i) = \{\phi_\lambda(X_i)\tilde{w}_\lambda : \psi_\lambda(X_i)\tilde{w}_\lambda = e_{i+1}(X_i), c_\lambda \in A_C\}$ be the a-priori synthesis of a basis, $\phi_\lambda \in A_C$, based on measurements at existing sample locations, X_i . Defining $\|\hat{f}_i(X_i) - \tilde{f}_{i+1}(c_\lambda, X_i)\| = \|\tilde{e}(c_\lambda, X_i)\|$ as a measure of expected reduction in approximation error due to the a-priori introduction of basis ϕ_λ , the selection criterion for MSAS algorithm based on greatest a-priori error correlation (MSAS-GC) is:

$$\text{selected node, } c_{i+1}^* = \arg \max_{c_\lambda \in A_C} \{\|\tilde{e}_\lambda(c_\lambda, X_i)\|\}.$$

As is discussed in the next section, MSAS-GR simulates a depth first trajectory through the scale-space whereas MSAS-GC simulates a breadth-first trajectory.

E. MSAS Illustration

Some properties of MSAS algorithms are illustrated using the following test functions:

$$\begin{aligned} f(x) &= -1 + (1 - x + 19x^2)e^{-x^2} + 2 \sin(6\pi x)e^{-(x-0.73)^2} \\ f(x, y) &= 2e^{-1.25(x-1.2)^2 - 10y^2} + \text{erf}(x - .3) + \text{erf} \\ f(x, y, z) &= x^2 + y^2 + e^{-z^2}. \end{aligned}$$

1) *Sample Distribution*: To illustrate the sample distribution generated by MSAS, adaptive sampling of the 1D test function is considered. A maximum sampling resolution of $(b - a)/2^7$ is assumed. Thus $M=7$ and the maximum number of samples is $2^7 + 1$. The basis trees generated after 50 iterations of MSAS-GR and MSAS-GC are shown in Figures 4a and 4b, respectively. For both algorithms, the total number of modeled nodes is 53, including three in the bootstrap iteration.

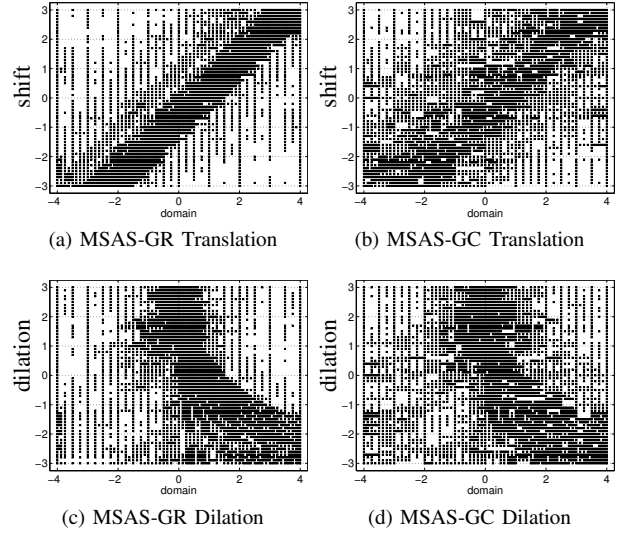


Fig. 6: Consistency of Sample Distribution (1D)

Figures 4a and 4b illustrate correspondence between the basis distribution and variation in the underlying function. Since the locations of un-modeled children (leaf nodes in the tree) of previously modeled nodes are the sample locations, the correspondence of bases distribution to variation extends to sample distribution. The sample distribution is shown in Figure 5 against the function. The two top rows corresponds to the sample distributions of MSAS-GR and MSAS-GC respectively. The bottom two rows represent ideal sample distributions generated using full knowledge of the underlying function a priori. Thus, MSAS achieves variation sensitive sample distribution through variation sensitive modeling of the underlying function.

MSAS-GC behaves as a breadth first algorithm and employs fewer level 7 bases than MSAS-GR. The differences in scale localization between the two MSAS basis trees arise because, in MSAS, variation in the underlying function is not explicitly characterized but derives only as a consequence of hierarchical error decomposition. Since the hierarchical decomposition of the error results in a model progresses from coarse scale basis to finer scale basis, in each iteration the residual error corresponds to variation in the underlying function. The scales at which a function is characterized is the scale at which the algorithm chooses to approximate the residual error at any intermediate iteration.

In MSAS, coarse-to-fine modeling is itself achieved by requiring only un-modeled children of previously modeled basis to enter the prospective basis pool. This ensures inter-generational continuity since a fine scale basis may only be added after either one of its parent bases at higher scale has been added to the model.

However, as a consequence of directly employing the sampling of the true approximation error rather than its estimate, MSAS achieves spatial localization of bases since regions of fine scale in the latter models are accounted for by

fine scale bases, even if at an apparent scale. The similarity in spatial localization of bases between the two algorithms is noted. For MSAS-GR, refinement of the 53 nodes yields 50 unmodeled child nodes. For MSAS-GC the corresponding number is 72.

2) *Consistency*: The spatial location of sample distribution may be characterized by considering the effect of translation and dilation of the underlying test function. The 1D test function $f(x)$ is independently subjected to translation and dilation through $f(2^d x - t)$, where $d \in \mathbb{R}$ is the dilation parameter and $t \in \Omega$ is the translation parameter. Figure 6 shows the consistency in sample distribution across various translations and dilations.

3) *Stability*: Since the internal hierarchical model, \hat{f}_k , employed by the algorithm represents a mechanism with which information content in sample distribution may be extracted, approximation error $e_{k+1} = f - \hat{f}_k$ may be used to characterize the relevance of a sample distribution. Since the approximation is based on localized sampling, the approximation error quantifies the appropriateness of the sample distribution.

In figure 7, average reconstruction error over various dilations and translations, and the variance in respective error averages are shown for all the test functions. Also shown is the reconstruction error due to MSAS Random Selection (MSAS-RS) algorithm. In each iteration of the MSAS algorithm, one basis is selected from the set of allowable candidate nodes with equal probability. The MSAS-RS error vs iteration curve shown is an average over 100 independent runs of the MSAS-RS algorithm for the test function.

As is seen in that figure, MSAS achieves a systematic reduction in the error through sequential addition of localized bases with local sampling. Further, the worst case performances of MSAS is atleast comparable to the best case performance of RS. Also, their average and the best case performance is significantly better than RS. Thus through sequential search and modeling, a sampling distribution that is representative of the underlying function is achieved.

4) *Higher Dimensions*: The sample distribution generated after $k = 50$ iterations of the MSAS-GR and MSAS-GC is shown overlayed on the test function in Figures 8a and 8b, respectively. The voronoi partition generated by the sample distribution is also shown to help visualize the spatial localization of the sample distribution. The sensitivity of the sample distribution to the variation in the underlying function is noted. In both algorithms, the number of bases introduced is 59, including 9 in the initial layer. The number of samples in MSAS-GR is 295, while the number of samples in MSAS-GC is 358. This indicates that, as in 1D, MSAS-GR behaves as a depth first algorithm and tends to select the middle child of a modeled node before either the left and right child.

In general, the characteristics of the MSAS algorithm in 1D are carried over to higher dimensions. The performance of the algorithm is invariant to translations and dilations of the underlying function, and the internal estimates of the algorithm converge to the original function. The resulting model is variation sensitive, as is the sample distribution.

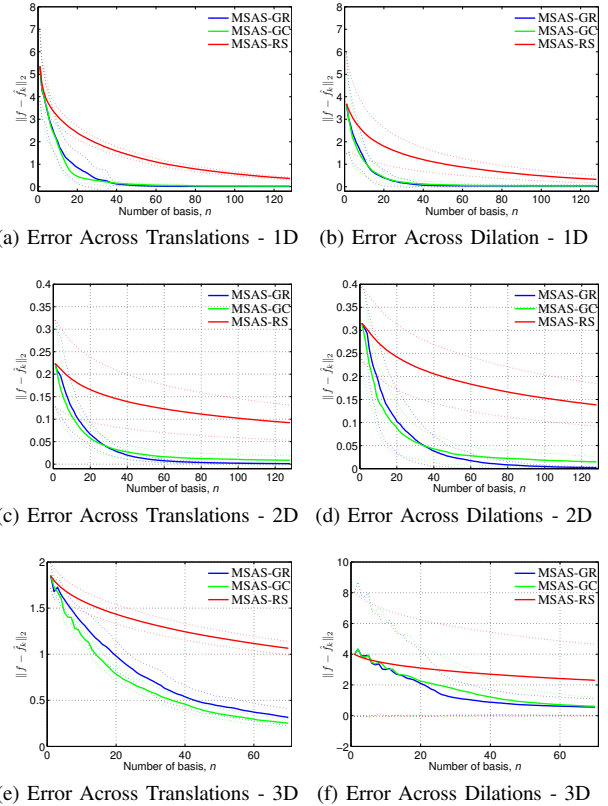


Fig. 7: MSAS Stability ($k = 50$)




Fig. 8: 2D Sample Distribution ($k = 50$)

Further, the relative performance of the algorithms is also retained. MSAS-GR behaves as a depth first algorithm while MSAS-GC behaves as a breadth first algorithm. However, in higher dimensions, the effect of modeling and sampling according to the structure of the grid becomes apparent. Due to the rigidity of the grid, while the bases and the sample distribution are sensitive to the variations of the underlying function, they are not aligned to the geometry of the underlying function. Instead, the sample distribution occurs as blocks aligned to the axis of analysis. This rigidity is active in 1D, however its effect is masked due to low dimensionality.

5) *Sampling Economy*: Variation sensitive sampling provided by MSAS provide significant benefits over uniform sampling. Figure 9 shows the ratio of the reconstruction




Fig. 9: MSAS Vs. Uniform Sampling 3D

error due MSAS-GC to reconstruction error due to uniform sampling as a function of sample rate for the 3D example. For a given sample rate $|X_i^{GC}| = n$ of MSAS-GC, the uniform sampling consists of the nearest square: $|X_i^{US}| = \lceil \sqrt{n} \rceil$. Since, in general $X_i^{US} \not\subseteq X_{i+1}^{US}$, the error due to uniform sampling is not guaranteed decrease monotonically, and the ratio curve is locally erratic. However, at each sample rate, MSAS offers significant economy in number of samples required for the same quality of reconstruction when compared to uniform sampling.

V. RESOURCE CONSTRAINTS

A typical sampling protocol at a given sample site involves sensor deployment, measurement and retrieval steps. With $\tau_m > 0$ as the time it takes to take a measurement at a sample location, the total “mission” time, T_s , may be written as: $T_s = L_s/v + k\tau_m \leq \Delta t$ where L_s is the total distance travelled by the vehicle in the mission, and v is the average speed of the vehicle during the mission. If the number of vehicles is fixed, then the process constraint $T_s \leq \Delta t$ imposes a tradeoff between coverage and spatial resolution. The vehicle speed is bound by a physical constraint, $v \leq V$. Further, the energy is also limited: $E_s = (\alpha v^3 + \beta) T_s \leq E_{\max}$, where α is a constant describing the vehicle’s hydrodynamic properties, and β is the hotel load. Thus the number of samples taken, k , constrains T_s since it increases hotel load and reduces energy available for transit, apart from taking up time τ_m . Increasing spatial resolution of measurements increases the time taken to sample a given area. In order to satisfy the temporal resolution imposed by the process constraints, the time taken to cover the area has to be reduced, which in turn implies reduced spatial resolution. If both spatial and temporal resolution have to be satisfied then the sample area has to be reduced. This fundamental tradeoff governs the decision about the appropriate time and space resolution and area to be covered.

Energy constraints may be introduced into MSAS by employing the augmented selection criteria: selected node $c^* = \arg \max_{c_\lambda \in A_C} (1-\theta)I(A_c) + \theta e^{-\beta \|c_\lambda - c\|^2}$, where $\beta > 0$ is a distance penalty parameter. Average tour distances and errors for MSAS-GR over random translations and dilations of the 2D function are shown in Fig 10 as a function of the θ . The tradeoffs are as expected. The key advantage is the tradeoff of distance travelled against information gained.

The frontiers of these trade-offs may be improved if multiple robots are employed. However, then issues such




Fig. 10: MSAS Tradeoffs (2D)

as co-operative routing need be considered. One approach to this problem is to seek clusters within the given sample distribution. Vehicles may then be assigned to their nearest cluster. A Hamiltonian path (Fig 1) may then be found within each cluster. This is the scope of future extension of this work.

ACKNOWLEDGMENTS

This work has been conducted in conjunction with the Center for Automation Technology at Rensselaer Polytechnic Institute and the Autonomous Undersea Systems Institute, Lee, NH.

REFERENCES

- [1] JT Allen, DA Smeed, and AL Chadwick. Eddies and mixing at the Iceland-Faroes Front. *Deep-sea research. Part 1. Oceanographic research papers*, 41(1):51–79, 1994.
- [2] M. Babin, JJ Cullen, CS Roesler, PL Donaghay, GJ Doucette, M. Kahru, MR Lewis, CA Scholin, ME Sieracki, and HM Sosik. New Approaches and Technologies for Observing Harmful Algal Blooms. *Oceanography*, 18(2), 2005.
- [3] M.A. Batalin and G.S. Sukhatme. Coverage, Exploration and Deployment by a Mobile Robot and Communication Network. *Telecommunication Systems*, 26(2):181–196, 2004.
- [4] David A. Cohn, Zoubin Ghahramani, and Michael I. Jordan. Active learning with statistical models. *CoRR*, cs.AI/9603104, 1996.
- [5] V.V. Fedorov, W.J. Studden, and E.M. Klimko. *Theory of optimal experiments*. Academic Press New York, 1972.
- [6] S. Ferrari, M. Maggioni, and NA Borghese. Multiscale approximation with hierarchical radial basis functions networks. *Neural Networks, IEEE Transactions on*, 15(1):178–188, 2004.
- [7] Vadiraj Hombal, Arthur C. Sanderson, and Richard Blidberg. A non-parametric iterative algorithm for adaptive sampling and robotic vehicle path planning. *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, pages 217–222, Oct. 2006.
- [8] N.E. Leonard, D.A. Paley, F. Lekien, R. Sepulchre, D.M. Fratantoni, and R.E. Davis. Collective motion, sensor networks, and ocean sampling. *Proceedings of the IEEE*, 95(1):48–74, Jan. 2007.
- [9] D.J.C. MacKay. Information-Based Objective Functions for Active Data Selection. *Neural Computation*, 4(4):590–604, 1992.
- [10] T. Poggio and F. Girosi. Networks for approximation and learning. *Proceedings of the IEEE*, 78(9):1481–1497, Sep 1990.
- [11] D.O. Popa, A.C. Sanderson, Vadiraj Hombal, R.J. Komerska, S.S. Mupparapu, R. Blidberg, and S.G. Chappel. Optimal sampling using singular value decomposition of the parameter variance space. *Intelligent Robots and Systems, 2005. (IROS 2005). 2005 IEEE/RSJ International Conference on*, pages 3131–3136, Aug. 2005.
- [12] M. Rixen and J.M. Beckers. A synopticity test of a sampling pattern in the Albanian Sea. *Journal of Marine Systems*, 35(1-2):111–130, 2002.
- [13] A. Singh, R. Nowak, and P. Ramanathan. Active learning for adaptive mobile sensing networks. In *Proceedings of the fifth international conference on Information processing in sensor networks*, pages 60–68. ACM New York, NY, USA, 2006.
- [14] J.S. Willcox, J.G. Bellingham, Yanwu Zhang, and A.B. Baggeroer. Performance metrics for oceanographic surveys with autonomous underwater vehicles. *Oceanic Engineering, IEEE Journal of*, 26(4):711–725, Oct 2001.