

Advantageous Change of Support in Inverse Implementations of Statistical Differential Equation Models

Mevin B. Hooten * Martha J. Garlick* James A. Powell*

Abstract

Statistical models using partial differential equations (PDEs) to describe dynamically evolving natural systems are appearing in the scientific literature with some regularity in recent years. Often such studies seek to characterize the dynamics of temporal or spatio-temporal phenomena such as invasive species, consumer-resource interactions, community evolution, and resource selection. Specifically, in the spatial setting, data are often available at varying spatial and temporal scales. Additionally, the necessary numerical integration of a PDE may be computationally infeasible over the spatial support of interest. We present an approach to impose advantageous changes of support in statistical implementations of PDE models and demonstrate its utility through simulation.

Key Words: Approximation, Averaging, Hierarchical Bayesian Model, Multi-Scale Analysis, Partial Differential Equation, Spatio-Temporal Model

1. Introduction

With the exception of a few ground-breaking efforts (e.g., Hotelling 1927; Fisher 1937), the use of sophisticated ecological models involving derivatives in a rigorous statistical setting has only gained ground recently (e.g., Wikle 2001; Hooten and Wikle 2008; Cangelosi and Hooten 2009; Zheng et al. 2009). As these recent studies have shown, scientifically motivated statistical models can be specified and fit readily in a hierarchical modeling context and often with a computational Bayesian implementation (i.e., Markov Chain Monte Carlo, MCMC).

Such modeling efforts have shown great promise for making science-based inference on dynamic natural processes, though they are not without their challenges. Issues that commonly arise include: the chosen form of stochasticity, numerical stability of the deterministic equations, and discrepancy between the measurements and process in terms of spatial and temporal support.

With regard to the change of support issue in particular, the effects of spatial support changes on statistical inference are well described in the literature (e.g., Gotway and Young 2002). In the case of up-scaling spatial support specifically, common methods generally involve some form of spatial averaging or integration of the variable of interest over space (Ferreira and Lee 2007). Despite numerous suggestions in the literature to the contrary, many forms of aggregation in spatio-temporal statistical modeling projects still rely on over-simplified and non-scientific spatial and/or temporal scaling methods (i.e., arithmetic averaging) without regard for the inherent properties or dynamic features of the process under study. When more sophisticated approaches to handling the scaling are adopted, they are often based on statistical or computational properties and less so on the mathematical underpinnings of the process being modeled (e.g., Ferreira et al. 2006). In fact, a popular method in the broader context of dimension reduction is to model the process on some lower-dimensional manifold through the use of a transformation

*Department of Mathematics and Statistics, Utah State University, Logan, UT 84322

involving a set of basis functions (e.g., Royle and Wikle 2005; Hooten and Wikle 2007).

In what follows, we present a general approach to change of support that can be advantageous when utilizing partial differential equation (PDE) models in a statistical framework. The methodology we present reconciles the statistical and mathematical forms of multi-scale analysis and can be beneficial, both in terms of dimension reduction, as well as consistency with the dynamics implied by the chosen mathematical model. A detailed presentation of the mathematical approach to multi-scale analysis can be found in the mathematical literature (e.g., Murray 1980; Holmes 1995; Okubo and Levin 2001), thus, here we focus on the use of multi-scale methods for approximating PDE solutions in a statistical context. We illustrate both the mathematical and statistical analyses through a specific example involving ecological diffusion and provide a simulation study and discussion of the advantages and potential applications of this approach.

1.1 Partial Differential Equation Example: Ecological Diffusion

To begin, we will focus our analysis on the specific example of ecological diffusion. To gain an understanding of how this form of mathematical model arises and how it may be useful, we derive it from first principles. Following Turchin (1998), we first consider an individual-based or Lagrangian understanding of animal movement. To illustrate the approach, we will use a one-dimensional spatial domain throughout, which easily generalizes to higher dimensions.

Suppose that, in a certain time interval, an animal can move left, right, or remain where it is with probabilities: $\phi_L(x, t)$, $\phi_R(x, t)$, and $\phi_N(x, t)$, respectively (where, $\phi_L(x, t) + \phi_R(x, t) + \phi_N(x, t) = 1$). Then the probability of the animal occupying location x at time t is:

$$\begin{aligned} p(x, t) &= \phi_L(x + \Delta x, t - \Delta t)p(x + \Delta x, t - \Delta t) \\ &\quad + \phi_R(x - \Delta x, t - \Delta t)p(x - \Delta x, t - \Delta t) \\ &\quad + \phi_N(x, t - \Delta t)p(x, t - \Delta t) \end{aligned} \quad (1)$$

where the Δ notation refers to changes in time and space. If we ultimately seek an Eulerian model on the probability of occupancy, $p(x, t)$, then we need to replace the Δ notation with differential notation. Turchin (1998) proceeds by expanding each of the probabilities in a Taylor series, truncating to remove higher order terms, and then substituting the truncated expansions back into (1). This yields a recurrence equation involving partial derivatives:

$$\begin{aligned} p &= (\phi_L + \phi_N + \phi_R)p - \Delta t(\phi_L + \phi_N + \phi_R) \frac{\partial p}{\partial t} - \Delta t p \frac{\partial}{\partial t} (\phi_L + \phi_N + \phi_R) \\ &\quad - \Delta x(\phi_R - \phi_L) \frac{\partial p}{\partial x} - \Delta x p \frac{\partial}{\partial x} (\phi_R - \phi_L) \\ &\quad + \frac{\Delta x^2}{2} (\phi_L + \phi_R) \frac{\partial^2 p}{\partial x^2} + \Delta x^2 \frac{\partial p}{\partial x} \frac{\partial}{\partial x} (\phi_L + \phi_R) + p \frac{\Delta x^2}{2} \frac{\partial^2}{\partial x^2} (\phi_L + \phi_R) \\ &\quad + \dots \end{aligned}$$

where we have defined $p \equiv p(x, t)$ to simplify the expressions. Combining like terms results in a PDE of the form:

$$\frac{\partial p}{\partial t} = -\frac{\partial}{\partial x}(\beta p) + \frac{\partial^2}{\partial x^2}(\mu p),$$

where, $\beta = \Delta x(\phi_R - \phi_L)/\Delta t$ and $\mu = \Delta x^2(\phi_R + \phi_L)/2\Delta t$. This resulting model is now Eulerian and known as the Fokker-Planck or Kolmogorov equation (Risken, 1989), and when thought of in an ecological setting, some number of total animals (N) can be considered by letting $u(x, t) \equiv Np(x, t)$. In this context, assuming for the moment that there is no advection component, we have the ecological diffusion equation:

$$\frac{\partial u}{\partial t} = \frac{\partial^2}{\partial x^2}(Du), \quad (2)$$

where, the process of interest is $u \equiv u(x, t)$, and $D \equiv D(x, t)$ represents the diffusion coefficients which can be thought of as animal motility. There are, of course, other ways to arrive at this equation (e.g., Turchin 1998), however, we feel that this perspective may be directly beneficial to those modeling spatio-temporal population dynamics, as the recent literature suggests. The properties of (2) are quite a bit different than those of plain or Fickian diffusions. The fundamental difference is that the diffusion coefficient is on the inside of the two spatial derivatives rather than between them (Fickian) or on the outside (plain). This results in a much less smooth process $u(x, t)$, and can be thought of as animals moving at varying rates through different environments on a landscape. In some areas, they may move slow, perhaps to forage, whereas in other areas they move fast, such as through an exposed terrain. The resulting behavior shows a congregative effect in areas of low motility (i.e., $D \downarrow$) and a dispersive effect in areas of high motility (i.e., $D \uparrow$). In fact, depending on the boundary conditions, the steady-state solution has u proportional to the inverse of D .

1.2 Statistical Implementation

Recent efforts in spatio-temporal statistical modeling (e.g., Wikle 2003; Hooten and Wikle 2008) suggest that we can place the above deterministic equation into a statistical context by assuming three things: first, that the dynamic process can be measured, and those observations are subject to sampling error; second, that the Eulerian model in (2) does not define the exact dynamics of the system, but rather, serves as a well-founded scientific motivation for parameterizing the dynamics; third, we seek to estimate the model parameters (e.g., D) given the observed data. Thus, we have two explicit sources of uncertainty in the model, observation uncertainty and process uncertainty. As discussed throughout the recent statistical literature on the subject, this situation is ideal for the use of a hierarchical model (e.g., Berliner 1996; Cressie et al. 2009). In this setting, we can think of the measurements as observed counts of animals, $N(x, t)$, at a set of locations and times. Then, by modeling the latent process $u(x, t)$ we can learn about the dynamics. Thus, using vector notation to denote finite sets of state variables and parameters (e.g., $\mathbf{u}(t) \equiv (u(1, t), \dots, u(x, t), \dots, u(n, t))'$ for n spatial locations of interest), let the statistical model be specified as,

$$\begin{aligned} N(x, t) &\sim [N(x, t)|u(x, t)] \quad , \quad \forall x, t \quad , \\ u(x, t) &\sim [u(x, t)|f(\mathbf{u}(t - \Delta t), \mathbf{D})] \quad , \quad \forall x, t \quad , \end{aligned}$$

where, the '[·]' notation refers to a probability distribution and the ecological diffusion model (2) is represented by a discretized approximation $f(\mathbf{u}(t - \Delta t), \mathbf{D})$. Typically the discretized approximation amounts to a difference equation, though more sophisticated solvers could be employed. Note that, in some situations, the specific form of f could facilitate estimation (i.e., if linear with Gaussian error, then

conjugacy is a byproduct). Then, assuming that a finite set of unique coefficients can be modeled by $\mathbf{D} \sim [\mathbf{D}]$, in a Bayesian framework we seek to find the conditional distribution (i.e., posterior) of the parameters \mathbf{D} and latent state variables $\mathbf{u}(t)$ given the data $\mathbf{N}(t)$:

$$\begin{aligned} \{ \{ \mathbf{u}(t) \}, \mathbf{D} | \{ \mathbf{N}(t) \} \} &\propto [\text{Data} | \text{Process}] [\text{Process} | \text{Parameters}] [\text{Parameters}] \\ &\propto \prod_t [\mathbf{N}(t) | \mathbf{u}(t)] \prod_t [\mathbf{u}(t) | f(\mathbf{u}(t - \Delta t), \mathbf{D})] [\mathbf{D}] \quad , \end{aligned} \quad (3)$$

where, for the time being, a set of initial and boundary conditions are assumed. In the generalized case, these conditions could be (and have been) further modeled.

This framework has proven to be a powerful means for directly incorporating scientific information into the modeling process and provides an intuitive way to account for observational and process uncertainty while estimating the parameters that control the dynamics. In principle, this statistical framework could be employed for any differential equation model if sufficient data were available to provide the learning. In practice, several complications can arise when fitting such models. First, if the model is fit using MCMC, the algorithm will require iterative evaluation of the discretized forward model f ; this can be very computationally demanding if the discretization is spatially and/or temporally fine. Additionally, there is often a discrepancy between the support of the measurements and that of the process but inference is desired at the finer of the two scales. In the specific situation of a spatially heterogeneous set of diffusion coefficients (i.e., $D(x)$), the dimension of the parameter space is potentially huge. Finally, there is the issue of stability of the forward model (f) itself and the fact that ecological diffusion is inherently less numerically stable than simpler forms (e.g., Fickian and plain diffusion). In the following section, we take an analytical approach to change of support in PDE modeling by introducing additional temporal and spatial scales in what is referred to as “homogenization” in the PDE literature (Holmes 1994; Pavliotis and Stuart 2008). We will see that such analyses can lead to a choice of statistical quantities that are consistent with the mathematics while inducing an advantageous change of support in the inverse implementation of the model. The result is a faster, more robust algorithm for fitting PDE models to data in a rigorous statistical framework.

2. Homogenization with Ecological Diffusion

In what follows, we use a compact mathematical notation for partial derivatives where subscripts on variables pertain to partial derivatives; for example, $u_t \equiv \frac{\partial u}{\partial t}$ and $\partial_x^2 u \equiv \frac{\partial^2 u}{\partial x \partial x}$. Now, focusing on the PDE in (2), $u_t = \partial_x^2(Du)$, we introduce two new scales, a fast temporal scale (τ) and a small spatial scale (y). Thus, the spatially varying diffusion coefficient can be written as: $D = D(x, x/\epsilon) = D(x, y)$ and similarly $u = u(x, y, t, t/\epsilon^2) = u(x, y, t, \tau)$, where the (small) variable ϵ represents the ratio of the scale at which D varies most rapidly and the scale on which inference is desired. To accommodate long and short scales, the derivatives are transformed (following Holmes 1996 and Powell and Zimmerman 2004) where, for small ϵ , $\partial_x \rightarrow \partial_x + \frac{1}{\epsilon} \partial_y$ and $\partial_t \rightarrow \partial_t + \frac{1}{\epsilon^2} \partial_\tau$.

Conventionally, in perturbation theory, a power expansion is employed so that one can break a difficult problem up into a series of solvable problems: $u = \epsilon^0 u_0 + \epsilon^1 u_1 + \epsilon^2 u_2 + \dots$. In this case, we substitute this expansion, along with both scales, into $u_t = \partial_x^2(Du)$ to yield:

$$(\partial_\tau + \epsilon^2 \partial_t)u = (\partial_y^2 + 2\epsilon \partial_{xy} + \epsilon^2 \partial_x^2)(D(x, y)u) \quad (4)$$

Now, we need only to solve each order of (4) sequentially. In doing so, we first concentrate on $\mathcal{O}(\epsilon^0)$: $\partial_\tau u_0 = \partial_y^2(D(x, y)u_0)$. Note that this equation involves a partial derivative with respect to the fast time scale τ , so it decays rapidly to the steady state solution: $\partial_y^2(Du_0) = 0$. Integrating once yields, $\partial_y(Du_0) = c_0(x, t)$ and then integrating again we have: $Du_0 = c_1(x, t) + yc_0(x, t)$. This latter term, $yc_0(x, t)$, grows unbounded on the short scale (termed “secular” in the perturbation theory literature) which is not consistent with the expected PDE behavior. Consequently, we set $c_0 = 0$ to yield $u_0 = c_1(x, t)/D(x, y)$. Note that this latter result yields densities that are proportional to D^{-1} , so that the population aggregates in regions of lower motility.

Next, we focus on the problem in $\mathcal{O}(\epsilon^1)$. In this order, one can show that (4) simplifies to: $\partial_\tau u_1 = \partial_y^2(D(x, y)u_1)$. This equation also decays to the steady state solution rapidly and the solution to this problem, $\partial_y^2(D(x, y)u_1) = 0$, is similar to that of $\mathcal{O}(\epsilon^0)$: $u_1 = b_1(x, t)/D(x, y)$. As we will see in the next result, this term will not contribute to the ultimate solution so we can set $u_1 = 0$.

The $\mathcal{O}(\epsilon^2)$ solution now becomes the critical part of the homogenization. From (4) we now have: $\partial_t u_0 + \partial_\tau u_2 = \partial_y^2(D(x, y)u_2) + 2\partial_{xy}(D(x, y)u_1) + \partial_x^2(D(x, y)u_0)$. Notice that the middle term on the right hand side goes away because $u_1 = 0$. Since τ is a fast variable, all terms based on $\partial_\tau u_2$ must vanish to avoid secularities in τ ; in other words, we need only to find the steady state solution with respect to τ . Substituting our earlier result for $u_0 = c_1(x, t)/D(x, y)$ we have:

$$\partial_y^2(D(x, y)u_2) = \partial_t \left(\frac{c_1(x, t)}{D(x, y)} \right) - \partial_x^2 c_1(x, t) . \tag{5}$$

Integrating (5) once yields:

$$\begin{aligned} 2\partial_y(D(x, y)u_2) &= \int_{y_0}^y \left(\partial_t \left(\frac{c_1(x, t)}{D(x, s)} \right) ds - \partial_x^2 c_1(x, t) \right) + a_0(x, t) \\ &= a_0(x, t) + \partial_t c_1(x, t) \int_{y_0}^y \frac{1}{D(x, s)} ds - 2y\partial_x^2 c_1(x, t) . \end{aligned} \tag{6}$$

To ensure that (5) is non-secular, we require:

$$\lim_{y \rightarrow \infty} 2y \left[\partial_t c_1(x, t) \frac{1}{2y} \int_{y_0}^y \frac{1}{D(x, s)} ds - y\partial_x^2 c_1(x, t) \right] = 0 . \tag{7}$$

The condition in (7) can now be rewritten as:

$$\partial_x^2 c_1(x, t) = \partial_t c_1(x, t) \left(\lim_{y \rightarrow \infty} \frac{1}{2y} \int_{y_0}^y \frac{1}{D(x, s)} ds \right) . \tag{8}$$

For notational convenience, we define the statistical quantities in (8) as $\bar{D} \equiv \left[\lim_{y \rightarrow \infty} \frac{1}{2y} \int_{y_0}^y \frac{1}{D(x, s)} ds \right]^{-1}$ and $\bar{u} \equiv c_1$. Statisticians might recognize the former quantity \bar{D} as the harmonic mean over some compact spatial support, and the latter quantity (\bar{u}) then represents a form of potential. Finally, we rewrite (8) as:

$$\frac{\partial \bar{u}}{\partial t} = \bar{D} \frac{\partial^2}{\partial x^2} \bar{u} . \tag{9}$$

Notice that this is a plain diffusion PDE, with the coefficient \bar{D} on the outside of the two partial derivatives rather than the inside, as in the ecological diffusion case (2).

The implications of this result are that our original ecological diffusion model can be well-approximated by a plain diffusion model on coarser spatial-temporal support where the new coefficient is the harmonic mean of the old coefficient over the change of support. The above derivation shows that the specific form of averaging employed here is mathematically consistent with PDE behavior and, since $u_0(x, t) = \frac{\bar{u}(x, t)}{D(x, y)}$, we can recover the approximate small scale process. Also, the practical benefit to using the harmonic mean in this situation is that it will be a smaller quantity than both the geometric and arithmetic means and is well-suited for averaging rates (e.g., diffusion coefficients).

From a deterministic perspective, the result derived above allows one to solve the PDE with a spatially varying diffusion coefficient over large domains with an efficient computational algorithm. The averaging alone can dramatically reduce calculations, however the homogenized equation (9) also describes the dynamics of a smooth potential field and this increases computational stability. One can then use a coarser discretization in the numerical solver for the forward problem. However, from a statistical perspective, we need to solve the inverse problem. That is, having observed $u(x, t)$ (or some discrete representation of it), we seek to learn about the parameters controlling the dynamics of the system (i.e., $D(x)$). If a Monte Carlo based estimation method (e.g., MCMC or importance sampling) were used, the above result would be absolutely necessary if inference were desired over large spatial domains due to the required iterative evaluation of a PDE solver.

In terms of implementation, one can proceed as described in Section 1.2 where, depending on the type of observations that are made, a suitable data model can be chosen along with an error distribution for the process model and prior for the parameters. To make the appropriate homogenization transformations, we can treat the harmonic mean as an operator (i.e., $\bar{\mathbf{D}} \equiv \bar{\mathbf{D}}(\mathbf{D})$) in the following way:

$$\begin{aligned} N(x, t) &\sim [N(x, t) | \bar{u}(x, t)] \quad , \quad \forall x, t \\ \bar{u}(x, t) &\sim [\bar{u}(x, t) | f_h(\bar{\mathbf{u}}(t - \Delta t), \bar{\mathbf{D}}(\mathbf{D}))] \quad , \quad \forall x, t \\ \mathbf{D} &\sim [\mathbf{D}] \quad , \end{aligned}$$

where, the function f_h represents the plain diffusion solver as a difference equation.

3. Simulation: 1-D Ecological Diffusion with a Finite Set of Distinct Environments

In order to illustrate the methods presented above, we consider a simulation of animal movement in a 1-D spatial setting with three different land types. The idea would be that animal motility is homogeneous within each distinct land type, but the landscape itself could be made up of irregular patterns involving those land types. This is a relatively realistic scenario where covariate information may be available for a certain spatial domain (e.g., remotely sensed data) and animals are counted at a set of locations over that landscape.

Suppose we code the distinct land types using dummy variables in an $n \times p$ design matrix \mathbf{X} , then we could write the spatial field of diffusion coefficients for our domain of interest as $\mathbf{D} = \mathbf{X}\mathbf{d}$, where \mathbf{d} is a vector of dimension p ($p \ll n$) containing the unique coefficients. As a simulation landscape, Figure 1 illustrates the irregularity of the land types and their associated diffusions. In this case, though we evaluated numerous scenarios, the specific one presented here has $\mathbf{d} = (0.5, 1, 1.5)'$. For comparison, we have also shown the harmonic mean \bar{D} , averaged over one twenty-fifth of the spatial domain or two units in the discretized case.




Figure 3: Forward simulation of $\bar{u}(x, t)$ via plain diffusion, shown on the same spatial domain. Time points are again indicated by intensity of gray shading.




Figure 4: Approximation of $u(x, t)$ over time as recovered from $\bar{u}(x, t)$ and \mathbf{D} .

simulated process by sampling $N(x, t) \sim \text{Pois}(u(x, t))$ for some finite set of locations x and times t . These observations will serve as simulated data and can be visualized in the same manner as the underlying dynamical processes (Figure 5). Note that there is substantially more noise in the observations than in the underlying process.




Figure 5: Simulated data, $N(x, t) \sim \text{Pois}(u(x, t))$, used to illustrate statistical implementation.

For clarity in this example and because we wish to compare the homogenized statistical model with the non-homogenized model in terms of parameter estimation, we will consider a simplified Bayesian model where the process stage is integrated over, and thus, collapsed into the likelihood. In both the homogenized and non-homogenized cases, we use a vague Gaussian prior on the diffusion coefficients $\mathbf{d} \sim N(\mathbf{1}, 1000^2 \mathbf{I})$. The likelihoods for the original PDE and homogenized PDE are then

conditionally independent Poissons: $\mathbf{N}(t) \sim \text{Pois}(f(\mathbf{u}(t - \Delta_t), \bar{\mathbf{D}}(\mathbf{d})))$ and $\mathbf{N}(t) \sim \text{Pois}(f_h(\bar{\mathbf{u}}(t - \Delta_t), \bar{\mathbf{D}}(\mathbf{d}))), \forall t$, respectively.

Conditioning on the initial state and no-flux boundary conditions, we then seek the posterior distribution for the diffusion coefficients from each model, both of which are conditional distributions of the form: $[\mathbf{d}|\{\mathbf{N}(t), \forall t\}]$. Due to the nonlinearities in the PDE solver as well as the specific forms of the likelihood and prior, the posterior distribution will not be conjugate, and thus, not analytically tractable. Following previous approaches to implementing such models (e.g., Wikle 2001), an MCMC algorithm can easily be constructed to sample from the posterior using Metropolis-Hastings updates for the diffusion coefficients \mathbf{d} .

For this simulation study, we obtained 10,000 MCMC samples and discarded the first 2,000 as the burn-in. The resulting chains for each model were converged, though we noticed better mixing in the case of the homogenized PDE model. We attributed this to the improved stability of the plain diffusion solver as compared with the ecological diffusion solver. Figure 6 illustrates the marginal posterior distributions for each of the coefficients with the original PDE results in black, the homogenized in red, and the truth shown as green vertical lines. Notice that each model captured the truth well and resulted in a high degree of posterior similarity, with the homogenized PDE posterior exhibiting a slightly reduced precision.




Figure 6: Marginal posterior distributions for each of the three diffusion (motility) coefficients. The results of the original PDE model fit are shown in black whereas the homogenized PDE appear in red and the “truth” is represented as vertical green lines.

In terms of computational performance, the homogenized PDE model took only 40% of the time it took the original PDE model to be fit. Clearly, the gain in computing efficiency is increased as the extent of averaging increases. That is, in this case, we only averaged over one twenty-fifth of the spatial domain. Through simulation we also evaluated the change in compute time as a function of both the change of support for averaging and the magnitude of the diffusion coefficients (i.e., $\max(\mathbf{D})$). The latter is important for the stability of the numerical PDE solver being employed; in this case, the solver is a centered difference equation. Figure 7 illustrates the empirical changes as a ratio of homogenized model computing time to original model computing time. Though some inconsistencies exist due to the empirical nature of the simulations, we can see that the computational savings generally improve with larger diffusion coefficients and larger changes in support. Notice however, that the improvement is asymptotic in both cases and for this simple 1-D ecological diffusion example, the best we can do with homogenization is

an algorithm that is 5 times faster. In the 2-D spatial setting, we might expect to improve on that by a power of two.




Figure 7: Computational improvements measured as the ratio of homogenized PDE to original PDE compute time (y-axis). The magnitude of the diffusion coefficients is shown on the x-axis and increases in change of support are shown in terms of gray shading intensity; larger scale averaging corresponds to darker shades.

4. Conclusion

In summary, by taking a mathematical approach to change of support in PDE models, we can obtain similar statistical inference on parameters controlling the dynamics of the process with a fraction of the computational effort. Though traditionally employed in engineering settings, the specific type of second order perturbation theory known as homogenization has both a statistical interpretation and utility. Homogenization allows for an accurate and efficient implementation of PDE models while still accommodating small scale structure in the underlying environment and is harmonious with existing multiscale statistical methods. In the approximation procedure for ecological diffusion, a natural, but uncommon, statistical quantity arises as a means for changing the support (i.e., the harmonic mean) and as a byproduct, we are left with a smoother dynamical system that can be implemented on a coarser spatial scale. Through simulation, we have demonstrated the use and effectiveness of this method and illustrated the potential savings in computation.

It is important to note that perturbation theory is applicable to most differential equations. In fact, in the epidemiological/ecological setting, we have found that a similar approach yields advantageous results in the implementation of 2-D ecological diffusion PDEs with growth terms as well as more complicated coupled PDE models (Garlick et al. 2009). Other applications of the methodology presented here include modeling of spatio-temporal atmospheric and environmental processes as well as epidemiological processes where the spread of disease is a primary concern.

REFERENCES

- Berliner, L.M. (1996), “Hierarchical Bayesian Time-Series Models,” In: *Maximum Entropy and Bayesian Methods*, 15-22. The Netherlands: Kluwer Academic Publishers.
- Cangelosi, A.R., and Hooten, M.B. (2009), “Models for Bounded Systems with Continuous Dynamics,” *Biometrics*, In Press.
- Cressie, N.A.C., Calder, C.A., Clark, J.S., Ver Hoef, J.M., and Wikle C.K. (2009), “Accounting for Uncertainty in Ecological Analysis: The Strengths and Limitations of Hierarchical Statistical Modeling,” *Ecological Applications*, 19, 553-570.
- Ferreira, M.A.R., Higdon, D., Lee, H.K.H, and West, M. (2005), “Multiscale and Hidden Resolution Time Series Models,” *Bayesian Analysis*, 1, 947-968.
- Ferreira, M.A.R. and Lee, H.K.H. (2007), *Multiscale Modeling, a Bayesian Perspective*, Springer.
- Fisher, R.A. (1937), “The Wave of Advance of Advantageous Genes,” *Annals of Eugenics*, 7, 355-369.
- Garlick, M.J., Powell, J.A., and Hooten, M.B. (2009), “Homogenization of Large-Scale Movement Models in Ecology with Application to the Spread of Chronic Wasting Disease in Ungulate Populations,” In Preparation.
- Gotway, C.A. and Young, L.J. (2002), “Combining Incompatible Spatial Data,” *Journal of the American Statistical Association*, 97, 632-648.
- Holmes, M.H. (1995), *Introduction to Perturbation Methods*, New York, Springer-Verlag.
- Hooten, M.B. and Wikle, C.K. (2007), “Shifts in the Spatio-Temporal Growth Dynamics of Short-leaf Pine,” *Environmental and Ecological Statistics*, 14, 207-227.
- Hooten, M.B. and Wikle, C.K. (2008), “A Hierarchical Bayesian Non-Linear Spatio-Temporal Model for the Spread of Invasive Species with Application to the Eurasian Collared-Dove,” *Environmental and Ecological Statistics*, 15, 59-70.
- Hotelling, H. (1927), “Differential Equations Subject to Error,” *Journal of the American Statistical Association*, 22, 283-314.
- Murray, J.D. (2002), *Mathematical Biology, 3rd Ed.* Spring-Verlag.
- Okubo, A. and Levin, S.A. (2001), *Diffusion and Ecological Problems: Modern Perspectives, 2nd Ed.*, Spring-Verlag.
- Pavliotis, G.A., and Stuart, A.M. (2008), *Multiscale Methods: Averaging and Homogenization*, New York, Springer-Verlag.
- Powell, J.A. and Zimmermann, N.E. (2004), “Multi-Scale Analysis of Seed Dispersal Contributes to the Resolution of Reid’s Paradox,” *Ecology*, 85, 490-506.
- Risken, H. (1989), *The Fokker-Planck Equation: Methods of Solution and Applications*, Springer.
- Royle, J.A. and Wikle, C.K. (2005), “Efficient Statistical Mapping of Avian Count Data,” *Environmental and Ecological Statistics*, 12, 225-243.
- Turchin, P. (1998), *Quantitative Analysis of Movement*, Sunderland, Massachusetts: Sinauer Associates, Inc. Publishers.
- Wikle, C.K., Milliff, R.F., Nychka, D., and Berliner, L.M. (2001), “Spatiotemporal Hierarchical Bayesian Modeling: Tropical Ocean Surface Winds,” *Journal of the American Statistical Association*, 96, 382-397.
- Wikle, C.K. (2003), “Hierarchical Bayesian Models for Predicting the Spread of Ecological Processes,” *Ecology*, 84, 1382-1394.
- Zheng, Y., Zhu, J., and Aukema, B. (2009), “Hierarchical Dynamic Modeling of Spatio-Temporal Binary Data,” In Review.