Precision Agric (2015) 16:275–296
DOI 10.1007/s11119-014-9375-4

A probabilistic Bayesian framework for progressively updating site-specific recommendations

Patrick G. Lawrence · Lisa J. Rew · Bruce D. Maxwell

Published online: 9 September 2014

Abstract The goal of this research was to create an agricultural adaptive management framework that enables the probabilistic optimization of N fertilizer to achieve maximized net returns under multiple uncertainties. These uncertainties come in the form of bioclimatic variables that drive crop yield, and economic variables that determine profitability. Taking advantage of variable rate application (VRA), spatial monitoring technologies, and historical datasets, we demonstrate a comprehensive spatiotemporal modeling approach that can achieve optimal efficiency for the producer under such uncertainties. The utility of VRA fertilizer research for producers is dependent upon a localized accurate understanding of crop responses under a range of possible climatic regimes. We propose an optimization framework that continuously updates by integrating annual on-site experiments, VRA prescriptions, crop prices received, input prices, and climatic conditions observed each year under a dryland spring wheat (*Triticum aestivum*) cropping system. The spatio-temporal Bayesian framework used to assimilate these data sources also enables calculation of the probabilities of economic returns and the risks associated with different VRA strategies. The results from our simulation experiments indicated that our framework can successfully arrive at optimum N management within 6–8 years using sequential Bayesian analysis, given complete uncertainty in water as a driver of crop yield. Once optimized, the spatial N management approach increased net returns by \$23–25 ha^{-1} over that of uniform N management. By identifying small-scale targeted treatments that can be merged with VRA prescriptions, our framework ensures continuous reductions in parameter uncertainty. Thus we have demonstrated a useful decision aid framework that can empower agricultural

P. G. Lawrence (✉) · L. J. Rew · B. D. Maxwell
Land Resources and Environmental Sciences Department, Montana State University, Bozeman, MT 59717-3120, USA
e-mail: patrick.lawrence@msu.montana.edu; pat_nomad@yahoo.com

L. J. Rew
e-mail: lrew@montana.edu

B. D. Maxwell
e-mail: bmax@montana.edu

producers with site-specific management that fully accounts for the range of possible conditions farmers must face.

Keywords Site-specific experimentation · Bayesian statistics · Input optimization · Simulation experiment · Dryland agriculture · Spatial variation

Introduction

Variable rate application (VRA) research within the last 15 years has focused on finding the optimal spatial arrangement of fertilizer for maximizing net return and yield (Thrikawala et al. 1999; Mamo et al. 2003; Anselin et al. 2004; Liu et al. 2006; Biermacher et al. 2009; Meyer-Aurich et al. 2010). While these efforts have advanced understanding of the factors driving variation in crop productivity, the predominant outcome has been a fertilizer prescription map that is implied to be the best management strategy for an unspecified number of subsequent years (e.g. Anselin et al. 2004). However, given the short temporal scale of the datasets that produce these prescriptions, there remains substantial uncertainty in the expected crop responses. Such uncertainty, often caused by climatic variability (Mamo et al. 2003; Lambert et al. 2006; Florin et al. 2009), could be reduced by incorporating information from additional years of data. That is, if each year's prescription and response were assimilated into the data model and used to generate the following year's treatments, then predictions would continually improve, and a degree of adaptability would be built into the Precision Agriculture (PA) system.

Single-year prescriptions

Most VRA research to date has used some version of a randomized complete block strip experiment where fertilizer was applied at different rates to areas of the field that had been stratified based on some prior knowledge of crop response (Anselin et al. 2004; Liu et al. 2006; Shahandeh et al. 2010). These studies were typically based on 1 year of data, however there are isolated examples where additional years were incorporated (Shahandeh et al. 2005; Lambert et al. 2006). Using some form of linear or quadratic model that incorporates site-specific (for example soil texture), and occasionally year-specific variables, a regression on yield is then performed, which feeds into a net return function (Koch et al. 2004). This net return function is then maximized by finding the optimal value of fertilizer (usually nitrogen—N) to apply, and a prescription map is generated for each location in space or for management zones (Khosla et al. 2008).

A principal limitation of creating prescription maps based on 1 year of data is that it ignores the largest source of yield-limiting variability: climate (especially in dry climate systems). For example, if a prescription map based on an anomalous wet year were to be implemented during drier growing seasons, there would likely be a build-up of nitrogen in all but the most moist and responsive areas of the field. Furthermore, the ability of additional years of data to improve prescription performance is ignored, and the producer is left with a static recommendation that may offer no improvement over uniform application. Such condition-specific responses may at first glance suggest development of prescriptions for multiple precipitation scenarios, but until climate can be more accurately predicted, such approaches will not be fruitful. Maximizing the economic efficiency and minimizing

pollution while accounting for spatiotemporal variation over multiple years would instead enable decision-making that is more robust to the full range of possible conditions.

Multi-year prescriptions: accounting for spatial and temporal variability

The primary hurdles for incorporating spatiotemporal variability and non-experimentally controlled fertilizer applications relate to technical difficulties and the availability of data. Despite the increasing ubiquity of PA equipment, many producers do not have a complete, organized, and consistent set of records that span multiple years. Without such temporally consistent data, it is impossible to begin accounting for climatic variation, and to assess trends over time. Yields and yield responses to N can vary drastically across years, particularly in dryland small grain systems, so incorporating multiple cropping cycles is extremely important for decreasing prediction error (Kravchenko et al. 2005; Sadler et al. 2007; Florin et al. 2009). Adding additional years of data could be accomplished by automating the process of annual data assimilation and model updating.

With multiple years of data that are spatially co-located, statistical dependencies arise between observations that are close in space and in time, increasing the difficulty of achieving unbiased estimates of model parameters. The primary approach to minimize the influence of spatial and temporal autocorrelation is through the use of within-group or between-group modeling structures, using either **G**-side (sometimes denoted as the matrix **D**) or **R**-side covariance matrices (Laird and Ware 1982; Robinson 1991). **G**-side covariance matrices are used in mixed models where each group is distributed as $N(0, \mathbf{G})$, independently of the other groups and the errors. **R**-side covariance matrices are used for repeated measures or spatial analysis where autocorrelation is a structural component of the errors, and are distributed as $N(0, \mathbf{R})$. Depending on the dataset and autocorrelation of the bioclimatic variables, either of these approaches may be appropriate to accomplish the goal of reducing bias in the parameters or response predictions.

The most obvious means of incorporating temporally dependent inter-annual climatic variation is to use model parameters that represent the prevailing meteorological conditions such as precipitation, temperature or solar radiation for each year of production. All yield observations within 1 year are likely to be temporally autocorrelated due to the unique climatic conditions that occur in one season, leading to a field-wide annual bias in yields. This autocorrelation can be addressed by inducing within-year correlation via mixed models (Thöle et al. 2013). If blocked experimental plots or strips were used in the experimental design, a crossed random-effects modeling structure (**G**-side) can account for the spatial autocorrelation in yield values caused by unobserved soil, topography or ecological factors, as long as the repetitions are sufficiently separated in space (within a field) to ensure independence. However, a crossed random-effects structure does not adequately deal with spatial autocorrelation of yields in situations where the treatments are continuous across a field, as would be expected in a real farm scenario where an N prescription map is being used (Fig. 1). In such a situation where each unit in the field (cell) is informing the statistical model, the block-specific random effects do not account for the spatial autocorrelation of nearby cells because those random effects themselves are assumed to be independent (not true in an **R-**side approach) regardless of the distance by which they are separated.

Although spatial variation often causes smaller differences in yield than temporal variation (Florin et al. 2009), it can be much more difficult to incorporate into a crop yield model. The complete characterization of variation in soil properties outside (even inside) strictly controlled experimental settings is nearly impossible, and it is very difficult to even

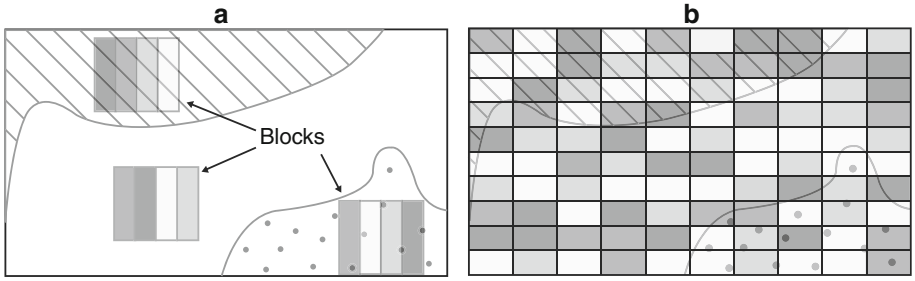

Fig. 1 The typical experimental layout of VRA (Bongiovanni et al. 2007), set against a background of edaphic spatial variation (*diagonal strikes* sandy, *dots* clayey, *empty* loam). In the randomized block design of **a**, there is correlation within blocks but not between, which can be easily dealt with using a random effect for each block (or fixed effects if the soil variation is of interest). In **b** the lack of discrete randomized blocks and the occurrence of cells on the transition zones between edaphic conditions makes the use of a mixed effects correlation structure inappropriate

know the scale of edaphic variation (Cambardella and Karlen 1999; Baxter et al. 2003; Kerry and Oliver 2003; Patzold et al. 2008). Measurements of apparent electrical conductivity (EC_a; mS m^{-1}) provide a potential way to efficiently collect continuous soils data, however the soil properties they measure (salinity, texture, and water content) are interrelated, making it difficult to interpret EC_a values (Corwin and Lesch 2003, 2005). Nevertheless, since these soil properties directly influence plant yield (Jung et al. 2005; King et al. 2005; Kühn et al. 2008), EC_a still may be adequate for characterizing edaphic variation if measured when soils are moist (Brevik et al. 2006). Other sources of spatial variation, such as topography and weed/pest pressure, can also be quantified, although biological organisms are difficult to measure in a spatially and temporally dense manner.

Another method to incorporate spatial variation, other than using random effects (**G**-side) for blocking factors, requires the use of a covariance matrix (**R**-side) that accounts for spatially autocorrelated errors (Lambert et al. 2006). By using an appropriate covariance matrix, the bias in the model parameters is reduced, although it can increase the difficulty of validating the fit and parameters chosen for the semivariogram. Inverse meta-modeling is another method that aims to use observed spatial variation in yield to inversely derive soil properties such as Available Water Capacity (AWC) (Florin et al. 2010), but if VRA has already been implemented on a field it may no longer be possible to separate inherent from management-induced variation.

A third method to deal with spatial autocorrelation of yield model residuals is Conditional Auto Regression (CAR) (Besag 1974; Jiang et al. 2009 for a PA example), which relies on the commonly derived lattice structure of aggregated PA data (Anselin et al. 2004). The advantage to working with the CAR model is that it is a computationally efficient way of managing fine-scale spatial statistical dependencies while allowing covariates to capture broader-scale trends (Lichstein et al. 2002). Using a simple linear regression, the baseline model for yield (Y) for each location i, omitting the CAR adjustment, is as follows:

$$Y_i \sim N\left(\mu, \sigma_e^2\right) \tag{1}$$

where every Y_i has the same variance, and the covariance between yield at location i and yield at another location j is modeled as zero (resulting in residual spatial autocorrelation if

present). In contrast, the CAR model adds a spatial random effect ϕ to the model for the mean:

$$Y_i \sim \left(x_i^{'}\beta + \phi_i, \sigma_e^2\right) \tag{2}$$

where x_i' represents a vector of covariates, β is the associated parameter values, and σ_e^2 represents the independent and identically distributed (i.i.d.) errors. The set of conditional distributions used to account for spatial autocorrelation among yield responses (ϕ_ι in Eq. 2) is specified by:

$$Y_i|Y_{j,j\neq i} \sim N\left(\sum\nolimits_j b_{ij}y_j, \tau^2\right) \tag{3}$$

where j represents the set of cells neighboring cell i. Y_i is thus characterized by a normal distribution function, with its mean conditional upon the average yield values y_j for neighboring cells. The b_{ij} are entries in the n x n symmetric matrix **B**, referred to as a spatial weights matrix, with all b_{ii} equal to 0, all b_{ij} adjacent to cell i equal to 1 (dependent observations), and all other b_{ij} equal to 0 (independent observations). Defined as such, this model for yield responses subsumes information from adjacent cells to arrive at parameter estimates for the focus cell in such a way that bias from spatial autocorrelation is reduced. If appropriate, the weights matrix may be modified to include more distant observations (second-order or higher). The variance–covariance matrix (not derived here) associated with the specification shown above is:

$$V = (\boldsymbol{I}_n - \boldsymbol{B})^{-1}\boldsymbol{M}$$

where $\boldsymbol{I}_n$ is the identity matrix, **B** is the aforementioned spatial weights matrix, and **M** is equal to $\sigma_e^2 I_n$. The complete specification for the CAR model also requires delineation of the prior distribution (prior distributions explained under the methods section) for ϕ_i, which is defined as follows:

$$\phi_i|\phi_{j\neq i} \sim N\left(\overline{\phi}_i, \frac{\tau_c^2}{m_i}\right), where\ \overline{\phi_i} = \frac{1}{m_i}\,\phi_j \sum\nolimits_{j\in\partial i}\phi_j \tag{5}$$

where ∂_i represents the set of neighbors surrounding cell i, and m_i is the number of these neighbors (Besag 1974; Besag et al. 1991). This implies that Y_i is conditioned both by the value of the explanatory variables but also by the adjacent yield values. As such, locations defined as neighbors have correlated random effects and non-neighboring locations have independent random effects.

Although alternate means of accounting for both spatial and temporal autocorrelation exist, none have been extensively promoted within the PA literature. Spatio-temporal CAR models (STCAR)s and dynamical spatio-temporal models (DSTM)s offer promising new approaches (Cressie and Wikle 2011), however they can be difficult to implement, require a larger set of temporal data than is likely to be available, and are computationally expensive. Despite this, they should be strongly considered as computational capacity increases and more extensive datasets become available.

Jiang et al. (2009) offered an advance in the PA research field that utilizes Bayesian statistics to implement a CAR model. The advantage to using Bayesian methods lies in the possibility of implementing hierarchical correlation structures and in achieving a concrete posterior probability rather than being forced to bootstrap parameter intervals obtained from frequentist statistics to obtain probability values. In addition, with Bayesian methods

it is possible to separate the process from the data and parameter models, and to allow for continuous parameter updating as more data become available (Gelman et al. 2004). The example of Jiang et al. (2009) relied on precipitation and temperature covariates to account for temporal autocorrelation, but residual temporal autocorrelation was not quantified. Regardless, the spatial autocorrelation appeared to be adequately modeled, which led to the elimination of any spatial patterns in the regression residuals.

Multi-year prescriptions: an adaptive approach

We propose an adaptive system that extends the approach of Jiang et al. (2009), to accomplish the goal of generating continually updating probabilistic prescriptions that improve over time. In doing so, we illustrate spatio-temporal hierarchical Bayesian modeling with simulated results based on dryland spring wheat yield data from Montana. The demonstrated framework and model estimates yield as a function of nitrogen, precipitation, and apparent soil electrical conductivity (EC_a).

To ensure a reasonable level of accuracy in the model, we sought to realistically capture physiological crop responses by utilizing a non-linear yield equation (Archontoulis and Miguez 2013). This yield equation was then integrated into a net return function for profitability analysis. From there, the model parameters and optimization could be annually updated to achieve our objectives of progressively improved crop yield forecasts, prescription maps and visualizations of the unexplored parameter space. Each of these components is used to produce and provide new experimental treatments for the producer. Together, the advances achieved by these objectives lay the groundwork for a more accurate and responsive PA system that is able to incorporate the multiple forms of uncertainty that farmers face, increasing adaptive capacity for an uncertain future.

Methods

To achieve our goal we employed Bayesian statistical theory as it relates to the ability for models to self-update. Briefly, the standard formulation for a Bayesian posterior distribution is as follows:

$$p(\theta|D) = \frac{f(D|\theta)p(\theta)}{p(D)} \tag{6}$$

where θ is the parameter(s) of interest, D is the observed data, and f is the likelihood function. Therefore, the probability of the parameter given the observed data is equal to the product of the likelihood and the prior belief in the distribution of θ, divided by the marginal probability of the data (equivalently represented as $\int f(D|\theta)p(\theta)d\theta$ in a continuous context). This suggests that the posterior probability of the parameter is a weighted combination of the observed data and the prior certainty about the value of the parameter. Knowledge about the prior distribution is solely based on knowledge from previous research or intuition, but in practice is usually uninformative when no prior data are available (as is the case in this study for the first year of observations). This leaves the posterior entirely determined by the observed data. However, if the prior itself is based upon a probability distribution that requires additional parameters for specification (hyperparameters λ), such as a Beta distribution that requires α and β parameters, then the posterior distribution is expanded as:

$$p(\theta|D) = \frac{p(D,\theta)}{p(D)} = \frac{f(D|\theta)p(\theta|\lambda)p(\lambda)}{p(D)} \frac{f(D|\theta)p(\theta|\lambda)p(\lambda)}{\iint f(D|\theta)p(\theta|\lambda)p(\lambda)d\theta d\lambda} \quad (7)$$

It is important to notice that the denominator in Eq. 7 is the marginal probability of the data when all of the parameters have been integrated over, and must equal one. In practice, this term is called the normalizing constant, and can be ignored if using computational (rather than analytical) methods to sample the posterior distribution.

Once 1 year of crop yield data has been collected, the resulting posterior can form the basis for a new prior distribution. In mathematical notation:

$$p(\theta|D', D) = \frac{f(D'|\theta)p(\theta|D)}{p(D')} \quad (8)$$

where D' is the current year's yield data, D is last year's yield data, and $p(D')$ is equivalent to $\int p(D\prime|\theta)p(\theta|D)d\theta$. The above equation simply demonstrates that during each time step, the old dataset can serve as the basis for the new prior, providing a platform for a continuously updating model that should become more precise after each time step, assuming that the process (yield) model is properly constructed.

To evaluate the posterior distribution, at any time step, it is possible to analytically derive the mathematical form of the posterior if the likelihood and prior functions are relatively simple and conjugate to each other. If this is not possible, then either Metropolis–Hastings or Gibbs (a more specific form of Metropolis–Hastings) sampling may be performed, whereby possible parameter values are proposed and either accepted or rejected in an iterative procedure. Over time, it can be mathematically shown that the parameter samples converge to the true joint posterior distribution (Gelman et al. 2004).

The non-linear bayesian car model

In the standard notation of linear regression, the yield model for cell i in year j (following the general logistic form in Archontoulis and Miguez 2013) used to address our objectives was as follows:

$$Yield_{ij} = \frac{\beta\max * precip_j}{1 + \exp\left(\beta_{shp} - \beta_1 * QuantN_{ij} - \beta_2 \times EC_{a,i} - \beta_3 * EC_{a,i} \times QuantN_{ij}\right)} + \phi_i + \varepsilon \quad (9)$$

where $\varepsilon \sim N\left(0, \sigma_e^2\right)$, and ϕ_i is the estimated spatial random effect associated with cell i, which is used to account for spatial autocorrelation (see Eq. 2). In this specification, the parameter β_{max} can be interpreted as the maximum amount of yield at the asymptote and β_{shp} can be interpreted as the shape parameter because it shifts the yield-N response curve to the left or right. EC_a represents the apparent electrical conductivity (mS m^{-1}) of the soil and serves as a proxy for soil properties that impact yield such as available water holding capacity. $QuantN_{ij}$ is the amount of nitrogen applied to cell i in year j in the form of urea (kg ha^{-1}). Details on the spatial random effect, the variance–covariance matrix, and the spatial weighting scheme are identical to the CAR model described in the introduction.

Although some potentially relevant factors such as temperature and soil N are omitted, this model sufficiently allows for the examination of the Bayesian framework as influenced by edaphic ($EC_{a,i}$; changes over space but not across years) and management ($QuantN_{ij}$) variables, plus an annually changing (precipitation) variable. The nonlinear term represents an asymptotic (logistic) response of yield to N application, which more closely resembles

the actual biological response function than a linear model (Archontoulis and Miguez 2013) except at very high toxic N rates, which a farmer is unlikely to ever apply. If this model was being used for the purpose of hypothesis testing, then it would be appropriate to transform the N variable to comply with the assumptions of linear regression. However, since the aim is simulation, for which unrealistic values such as negative numbers or infinitely increasing yields would skew optimizations, the more physiologically accurate non-linear approach is more appropriate.

From a Bayesian CAR perspective, the model was formulated as follows:

$$p\left(\beta_{\max}, \beta_{shp}, \beta_{1-3}, \sigma_e^2, \tau_c^2 | yield\right) = \frac{f\left(Yield|\beta_{\max}, \beta_{shap}, \beta_{1-5}, \sigma, \tau_c\right) p\left(\beta_{\max}, \beta_{shap}, \beta_{1-3}, \sigma_e^2, \tau_c^2\right)}{\int\int d\beta_{\max} d\beta_{shp} d\beta_{1-3} d\sigma_e d\tau_c * f\left(Yield|\beta_{\max}, \beta_{shp}, \beta_{1-3}, \sigma_e^2, \tau_c^2\right) \times p\left(\beta_{\max}, \beta_{shp}, \beta_{1-3}, \sigma_e^2, \tau_c^2\right)} \tag{10}$$

where the likelihood is modeled by a normal distribution, with $Y_{ij} \sim N\left(\mu, \sigma_e^2\right)$, and μ equal to nonlinear function (9) + ϕ_i, with a common spatial variance of τ_c^2 for the random effects. In our model (Fig. 2), we also specified a set of hyper-parameters on the β_1, β_2, β_3, β_{max}, β_{shp}, σ_e^2 and τ_c^2 parameters, which reflected our uncertainty in their prior distributions (Gelman et al. 2004; Jiang et al. 2009).

Integration into the net return-maximizing function

To enable economic analysis, the posterior distributions for the parameters of interest (β_{max}, β_{shp}, β_{1-3}) and for the nuisance parameter (ϕ_i) from the crop yield model were integrated into a net return function. In all years of the simulation, the precipitation amount changed and the crop price to be received at harvest was unknown. Therefore, uncertainty in both of these values was incorporated into the net return function following the approach of Anselin et al. (2004).

$$\begin{aligned} Net\,Return_{ij} = {} & Price_{crop,j} \\ & * E\left[\frac{\beta\max * precip_j}{1 + \exp\left(\beta_{shp} - \beta_1 * QuantN_{ij} - \beta_2 * EC_{a,i} - \beta_3 * EC_{a,i} * QuantN_{ij}\right)} + \phi_i\right] \\ & - PriceN_j * QuantN_{ij} - FC \end{aligned} \tag{11}$$

where E[] is the expected value of the yield function and ϕ_i the spatial random effect in cell i (conditional on the neighboring cell random effects), $PriceN_j$ is the price of N (dollars/kg) in the current year, $QuantN_{ij}$ is the quantity of N applied (kg/ha), and FC is other average fixed costs associated with crop management ($605.44/ha) (USDA 2012a). The difference between the previous formulation (Anselin et al. 2004) and our construction is that instead of using a fixed value for the crop price and expected values for the parameters, we used distributions:

$$\begin{aligned} Net\,return_{ij} = {} & p\left(price_{crop,j}|Histprice_{crop}\right) * p\left(Yield_{ij}|param\right) * p(param) \\ & - p\left(priceN_j|Hist\,\mathrm{Pr}\,iceN\right) * QuantN_{ij} - FC. \end{aligned} \tag{12}$$

Where net return is in $/ha, $p(Price_{crop,j}|HistPrice_{crop})$ is the posterior probability distribution of crop prices from an autoregressive time series model for a historical dataset of prices ($/kg), $p(Yld_{ij}|param)*p(param)$ is the posterior probability of yield values given the

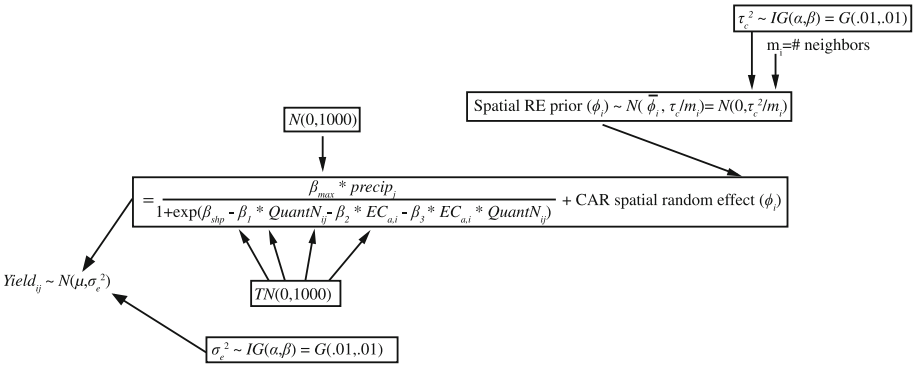

Fig. 2 Diagrammatic flow of the Bayesian CAR model with all distributions for the parameters and hyperparameters specified

uncertainty in the parameter values, the amount of applied N, and the cell-specific random intercepts, $p(priceN_j|HistPriceN)$ is the probability of current N prices given the historical distribution of nitrogen prices, $QuantN_{ij}$ is the quantity of N applied to cell i in year j, and FC represents the fixed costs.

Since the output of this function is a distribution on net return from which $QuantN_{ij}$ cannot be optimized, we instead performed Monte-Carlo simulations on the input distributions to obtain unique net return realizations. This was accomplished by drawing one set of possible parameters from the posterior parameter distributions, adjusting up or down to reflect the spatial effect of cell i and the unexplained variance, then optimizing N to obtain the maximum net return for cell i. This was performed 1,000 times in order to derive a distribution of the optimal values of N rate to apply to each cell. For a producer, the final parameter of interest would be the mean of the optimal N values for each cell, which integrated all of the uncertainty in the parameters including precipitation and the commodity price received (Fig. 3).

Annual updating

The net return-maximizing optimization process culminated in a site (cell)-specific N rate prescription map. The following year, the updating process began, inserting the posterior μ and σ values for each parameter as the new priors, and using new observations for the data. For each subsequent year, it was expected that the variance of each parameter would sequentially converge to the true variance, increasing certainty in the optimization prescriptions.

Fully exploring the parameter spaces

If a producer were to use the raw prescription map each year without modification, it is highly possible that some areas would receive exclusively large, or small, amounts of N. While this may be optimal based on the derived spatial crop responses to N, it could prevent the exploration of parameter combinations that might increase the net returns. For example, if only wet years (in an arid location where extra moisture is nearly always beneficial) had been observed, it might be assumed that under clayey soils the crop would

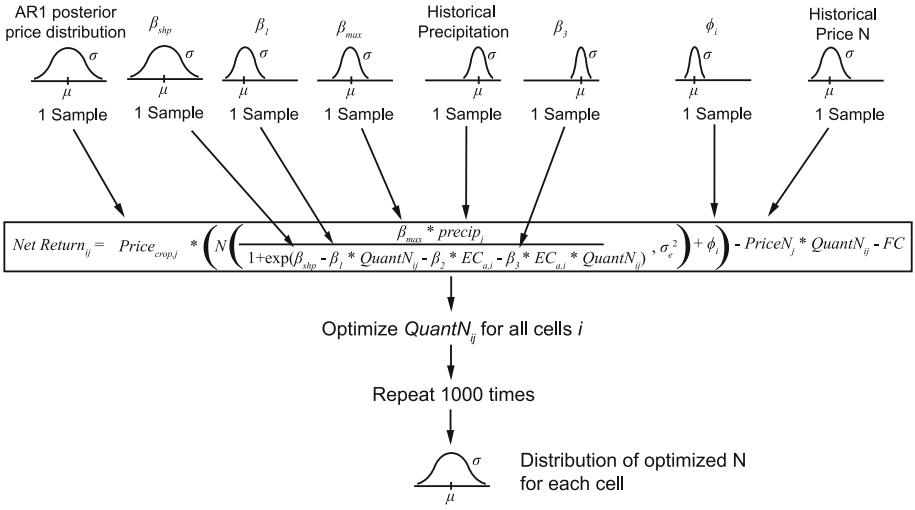

Fig. 3 The process by which crop yield parameter posterior distributions are sampled, N rate applied was optimized, and a distribution of optimized N was derived for each cell

respond favorably to high levels of N. However, that relationship might not hold true under dry years. This illustrates the importance of exploring the parameter space completely, at least during the initial years of PA implementation, in order to determine the crop responses under a full range of possible conditions. This exploration must be continued until enough years have been sampled to have high certainty in the distribution of conditions. Our framework accomplished this exploration through the use of annual N rate experiments on top of the optimized N prescriptions, and through visualizations of the N-EC-precipitation parameter space.

Simulation model implementation

The model as described above was implemented (Table 1) on a simulated 30-row by 30-column grid, with each grid cell representing a theoretical hectare. In practice, most fields are divided into smaller cell sizes that fit the size of the VRA fertilizing equipment (albeit with a similar number of total cells), however the use of one hectare cells provided easy interpretability for this example and was ideal for visualizing results. Mean values and variances for the input variables were chosen to be similar to those observed within multiple $\sim$50 ha fields (111.49°W, 47.68°N) growing non-irrigated spring wheat, located near Great Falls, Montana. We chose to use simulated rather than real data for the EC_a and ϕ_i variables in order to have complete knowledge of the structure of spatial variation. The spatially correlated EC_a Gaussian random field grid (Fig. 4) was generated within the R package RandomFields (R Core Team 2012; Schlather 2012) and was characterized by an exponential isotropic spatial covariance structure ($\sigma^2 = 640$, $\mu = 50$, range $= 50$, nugget $= 0$, scale $= 1$). Realized precipitation values for each year of the simulation were based on historical precipitation data (site: Sun River 4s) (National Climatic Data Center 2013), with the distribution centered at 26 cm and a standard deviation of 6.4 cm. Crop price received was based on a posterior distribution from a simple time series autoregressive lag 1 (AR1) model for first-differenced price data (1998–2012) obtained from the

Montana Wheat and Barley Committee (MWBC) (Montana Wheat and Barley Committee 2013) for current data; historical data obtained directly from the MWBC. The price uncertainty experienced by a farmer was approximated by obtaining the forecasted distribution 1 year (365 days) into the future. Fertilizer cost data are generally proprietary, therefore a normal distribution was used to approximate its uncertainty $N(\mu,\sigma) = N(\$0.55/\text{kg}, \$0.055\text{kg})$, and was based on annual fertilizer cost data available from the USDA (USDA 2012a). Finally, fixed costs for the producer were obtained from the USDA (USDA 2012b), and were \$605.44/ha omitting fertilizer costs, which were included in the model.

Initial conditions for the simulated updating process assumed that a farmer beginning to use PA technology would start with at least 1 year of yield monitor data under a uniform fertilizer application (140 kg/ha) before attempting to implement VRA. Following the first year of observing spatially variable yields, the field was stratified into three different yield classes with equal frequency (high, medium, low), within which different N rate treatments were applied. N rate treatments were selected to minimize influence on profitability (i.e. occupied small areas). These treatments as designed were three cells long within the direction of travel, which helped to ensure that the fertilizer spreader had adequate time to turn on, definitively spread the fertilizer, and turn off within the designated treatment area. The average farmer is unlikely to implement such a spatial experimental design themselves without substantial assistance, thus the implementation was automated as much as possible.

To calculate yields in the initial year and in subsequent iterations, Eq. (9) was applied using the parameter coefficients (Table 2). The β_{shp} parameter was fixed in order to eliminate its tendency to co-vary with the other exponential parameters (all parameters shifting up or down together, resulting in non-differentiable curves). It was based on generalized yield responses to N in Montana, where substantial yield gains from N additions typically occur between 0 and 80 kg/ha (Jackson 1998). Further variation was added to the yield for realism by drawing random values from a normal distribution (centered at zero and with a standard deviation of 270 kg/ha) then adding those values to each cell in each year. The value of the additional variance was based on observed residual variation from the aforementioned study site near Great Falls. The mean parameter values were taken as the "true" parameter values, which would later be estimated using the Bayesian MCMC process (Gelman et al. 2004).

The value for ϕ_i, the spatial random effect, was calculated from a multivariate normal distribution with a mean of zero and covariance matrix with σ's of 75 kg/ha (5,625 kg/ha σ^2) for neighboring cells, and 0 kg/ha for non-neighboring cells. These values were based on observed spatial autocorrelation from the previously mentioned field experiment (Fig. 4). Markov Chain Monte-Carlo (MCMC) simulations for the posterior distributions of the parameters were performed using the python programming language and the free python package pymc (Fonnesbeck et al. 2012). Previous implementations of CAR models have primarily been implemented with the software WinBUGS ("Windows version of Bayesian Updating using Gibbs Sampler", http://www.mrc-bsu.cam.ac.uk/bugs/winbugs/contents.shtml), however WinBUGS has not been updated since 2007, and we deemed it valuable to build our framework in an open source software package that was continuing to be developed and improved.

Prior distributions (as explained in the introduction) used for the coefficient parameters followed either normal or truncated normal distributions (Table 3) (Jiang et al. 2009). The truncated normal distributions were used in order to prevent the non-linear parameters from moving into unrealistic values in our system. The variances were set to be extremely large ($1e^{-12}$) in the first year in order to make the priors non-informative for both the normal and truncated normal distributions. If expert knowledge was available that could

Table 1 Iterative process for refining parameter estimates and optimizations

Year 0	1. Generate EC_a surface from exponential spatial correlation structure (Fig. 4) 2. Generate yield autocorrelation structure ϕ_i 3. Generate initial precipitation value $\sim N(\mu = 26.2$ cm, $\sigma = 6.4$ cm) 4. Apply N at uniform rate of 140 kg/ha = input N 5. Calculate initial yield surface with Eq. 9 (from input N, EC_a, precip, ϕ_i and additional random variation N(0,270 kg/ha))
Year 1	1. Stratify year 0 yield into 3 equal sized classes (high, medium, low) 2. Apply randomized block treatments within each class (0,60,120,180 kg/ha), uniform N elsewhere (140 kg/ha) = input N 3. Simulate a new precipitation value $\sim N(\mu = 26.2$ cm, $\sigma = 6.4$ cm) 4. Calculate year 1 yield with Eq. 9 (from input N, EC_a, precip, ϕ_i and additional random variation N(0,270 kg/ha)) 5. Run Bayesian CAR and extract posterior distributions 6. Optimize N for the next year based on samples from parameter posterior distributions
Years 2–6	1. Stratify year $t - 1$ yield into 3 equal sized classes (high, medium, low) 2. Apply randomized block treatments within each class (0,60,120,180 kg/ha) 3. Merge optimized N prescription and randomized treatment layers = input N 4. Simulate a new precipitation value $\sim N(\mu = 26.2$ cm, $\sigma = 6.4$ cm) 5. Calculate year t yield with Eq. 9 (from input N, EC_a, precip, ϕ_i and additional random variation N(0,270 kg/ha)) 6. Run Bayesian CAR with posteriors from year $t - 1$ as the year t priors and extract posterior distributions 7. Optimize N for the next year based on samples from parameter posterior distributions

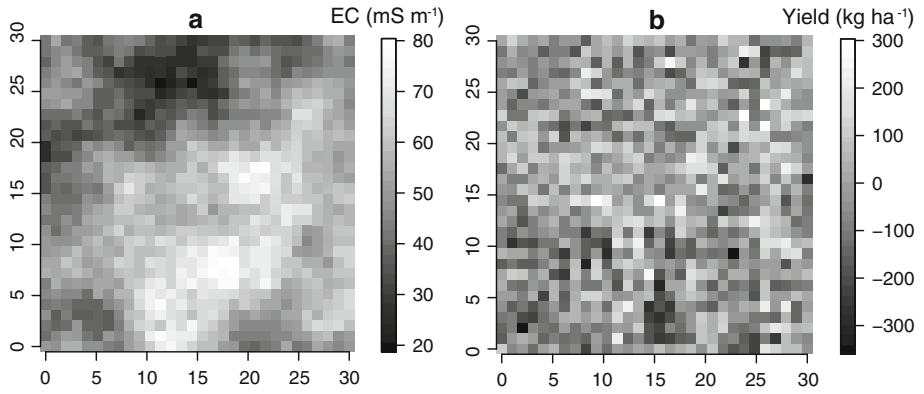

Fig. 4 EC_a surface with an exponential spatial autocorrelation structure (**a**), and surface of yield values (k/ha; ϕ_i) used to induce spatial autocorrelation in the response values (**b**)

Table 2 "True" parameters used to calculate yield within Eq. (9)

Parameter	β_{max}	β_{shp}	β_1	β_2	β_3	$\sigma_e = 1/\sqrt{\tau^2}$	$\sigma_s = \sqrt{\tau_c^2}/\sqrt{8}$
Value	137.8	4.8	0.02	.03	.0015	270	75

σ_e and σ_s are shown rather than τ^2 and τ_c^2 to enable the parameters to be interpreted on meaningful scales. Equivalent values for τ^2 and τ_c^2 -are .0000137 and 45 000 (parameterized as an inverse in python package pymc (Fonnesbeck et al. 2012) as .000022 (1/45 000))

direct the priors to be informative, then such knowledge could be incorporated initially, and would improve the convergence of the posterior distributions.

The prior distributions in the first year for the total model variance, σ_e^2 and the spatial variance (τ_c^2) were set to follow inverse-gamma distributions ($\sim IG(a, b)$) (Jiang et al. 2009), which were again specified to be non-informative (Gelman et al. 2004). In subsequent years, the priors were determined by the previous years' posterior distributions. During each year, the model was run for 100 000 iterations, using a burn-in period of 40 000 (wherein initial samples are discarded due to high autocorrelation) samples and a thinning rate of 20 in order to improve convergence and reduce autocorrelation between the samples. Convergence was confirmed through visual assessment of the parameter trace plots and autocorrelation plots.

Results and discussion

Convergence

Bayesian model validity depends on how well the specified model represents reality and the degree to which the posterior parameter distributions have converged. For this simulation model, distributions of the primary parameters of interest converged after 6 years of the simulation (excluding year zero), and the primary variance parameter σ_e required 8 years to converge (Fig. 5). The spatial variance parameter σ_s had not converged after 8 years, however it was trending towards its true value. Its lack of convergence was reasonable given all the other sources of confounding spatial variability, though it is

Table 3 Prior distributions for the coefficients (β), total variance (σ_e^2) and spatial variance (parameters τ_c^2) *TN* designates a truncated normal distribution

Parameter	Prior distribution with hyper-parameters	Hyper-parameter values	Prior distribution specification in pymc	Pymc hyper-parameter values
QuantN (β_1)	$TN(0,\sigma^2,a_N,b_N)$	*TN*(0.1, 1 E12, 0, 0.3)	*TN*(0.1, $1/\sigma^2$, a, b)	*TN*(0.1, 1 E-12, 0, 0.3)
EC_a (β_2)	$TN(0,\sigma^2,a_{EC},b_{EC})$	*TN*(0.1, 1 E12, 0, 0.5)	*TN*(0.1, $1/\sigma^2$, a, b)	*TN*(0.1, 1 E-12, 0, 0.5)
*QuantN*EC*$_a$ (β_3)	$TN(0,\sigma^2,a_{NEC},b_{NEC})$	*TN*(0.1, 1 E12, 0, 0.5)	*TN*(0.1, $1/\sigma^2$, a, b)	*TN*(0.1, 1 E-12, 0, 0.5)
β_{shp}	$TN(0,\sigma^2,a_{shp},b_{shp})$	*TN*(0.1, 1 E12, 2, 10)	*TN*(0.1, $1/\sigma^2$, a, b)	*TN*(0.1, 1 E-12, 2, 10)
Precip (β_{max})	$N(0,\sigma^2)$	*N*(0.0, 1 E12)	*N*(0.1, $1/\sigma^2$)	*N*(0.0, 1 E-12)
σ_e^2	$IG(\alpha_e, \beta_e)$	*IG*(0.01, 100)	*Gamma*(a_e, $1/b_e$)	*Gamma*(0.01, 0.01)
τ_c^2	$IG(\alpha_\tau, \beta_\tau)$	*IG*(0.01, 100)	*Gamma*(a_τ, $1/b_\tau$)	*Gamma*(0.01, 0.01)

possible that with an expanded CAR neighborhood size the parameter could be more accurately and quickly estimated. Additional simulations (not shown) using unique randomly drawn values of precipitation converged after 6–10 years, indicating that convergence can be achieved under different precipitation scenarios. In either case, multiple years of precipitation observations were required, reinforcing the need to collect and utilize multi-year data, and create N rate experiments over time.

Spatiotemporal variation

Residuals following year 6 of the simulations show minimal spatial pattern (Fig. 6), indicating that spatial autocorrelation was sufficiently accounted for (Moran's I = 0.01, p value for significant spatial autocorrelation = 0.48). Since temporal variation in yields was simulated using only one variable (precipitation), and that variable was included in the CAR model, temporal autocorrelation for each cell and for the field as a whole became insignificant once the model converged. This could be deduced from the lack of field-wide residual trends between years (after convergence), even under different precipitation conditions (Figs. 6, 7). If a longer set of years was observed, quantitative metrics rather than visual assessment could be used to assess the temporal autocorrelation. In a real world scenario with many possible drivers of temporal variation rather than only precipitation, it would be essential to assess residual temporal autocorrelation, which would give insight into the ability of temporal covariates to explain inter-annual variation. If the temporal covariates did not perform adequately, more climatic variables should be considered in the analysis.

The parameter space plots (Fig. 8) suggest that in addition to multiple years of precipitation observations, the experimental treatments applied in each year were crucial for achieving convergence. Within each year, the optimizations are visible as obvious clusters of points, whereas the fertilization treatments can be discerned by their correspondence to levels of 0, 60, 120 and 180 kg/ha of N. In the first 5 years of the simulation before convergence was achieved (years 5 and 6 omitted for clarity), the optimization chose N values that were near 0 kg/ha, whereas after convergence the optimization selected values clustered between 70 and 130 kg/ha. If only optimized N values were applied in each year (i.e. not applying rate experiments to explore the parameter space), convergence would be far less efficient or even impossible, especially for non-linear functions. The strategy of applying small experimental strip treatments, comprised of a range of N values in different yielding areas, helped speed convergence while eliminating the need for a farmer to devote their entire field to potentially non-profitable experimentation.

Small differences between the true parameters and the estimated parameters had negligible effects on the relationship between yield and N after year 6 (Fig. 9). With the exception of the spatial intercept, these parameters were identical for all areas in the field in each year. We chose to fix the value of the shape parameter due to its interdependence with β_1, β_2, and β_3, however given a sufficiently comprehensive dataset, it may be possible for the shape parameter to be estimated within the model. Furthermore, depending upon the purpose of the model, it may not matter whether the exponential parameters are interdependent if the realized Yield-N relationship remains unaffected. This behavior was observed during previous model runs when β_{shp} was not fixed.

Optimization

Maximizing net return and reducing the amount of unnecessary fertilizer applied requires knowledge of the underlying N-yield relationship and the response of net return to the full

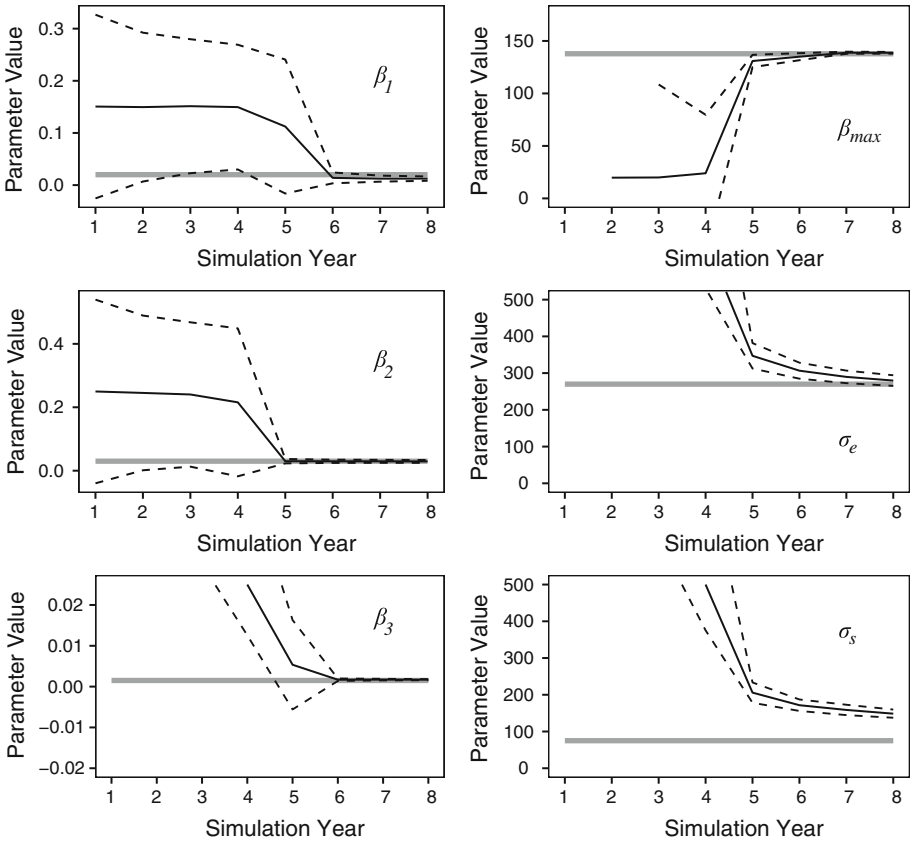

Fig. 5 Convergence of the posterior distributions in successive simulation years. *Solid lines* represent the posterior means, and *dotted lines* represent the ±2 standard deviations. The *gray lines* represent the true values of the parameters. β_1 = nitrogen, β_2 = EC_a, β_3 = nitrogen*EC_a, β_{max} = precipitation

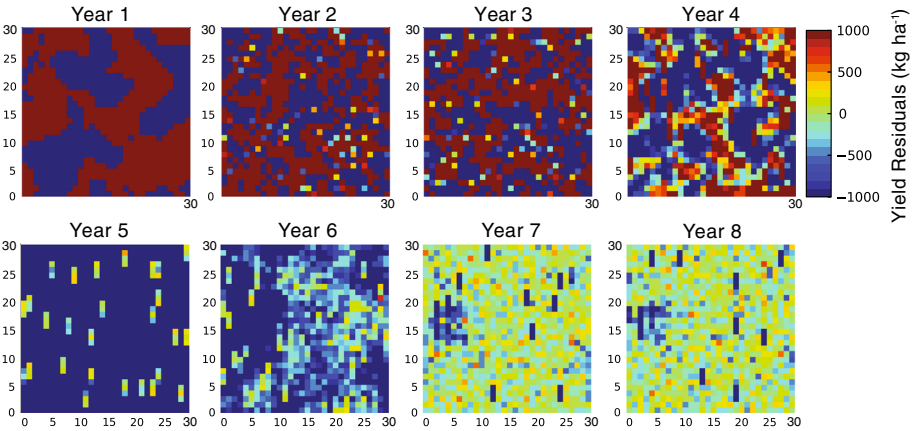

Fig. 6 Yield residuals for years 1–8 for the entire field, with convergence in year 6

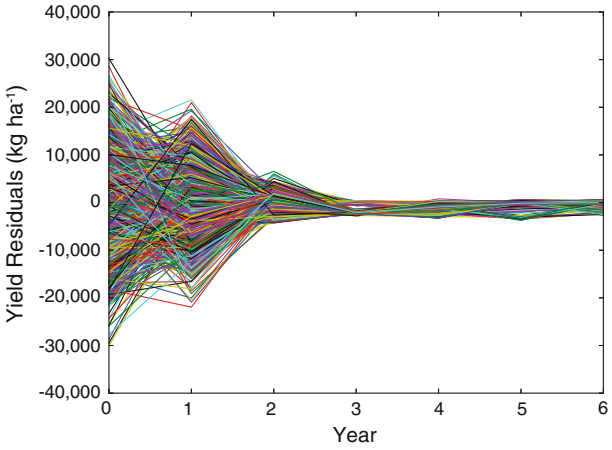

Fig. 7 Yield residuals for years 1–6 for each individual cell in the 900-cell field

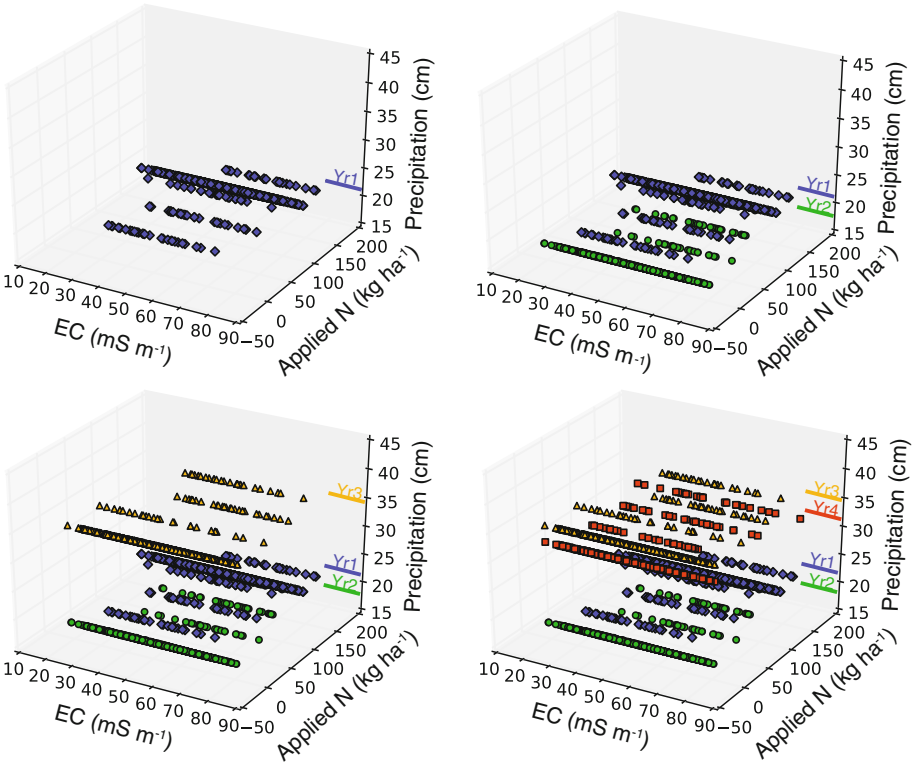

Fig. 8 Exploration of the first 4 years (year 0 with uniform N at 120 kg/ha and the next 3 years) of the precipitation-N-EC parameter space. *Blue diamonds* = year 1, *green circles* = year 2, *orange triangles* = year 3, *red squares* = year 4

range of bioclimatic and economic variability. Therefore, as a wider range of conditions are observed, the underlying crop parameters are better understood and the optimization is more efficient at maximizing net returns. Predictably, the optimizations for years 1–5 resulted in

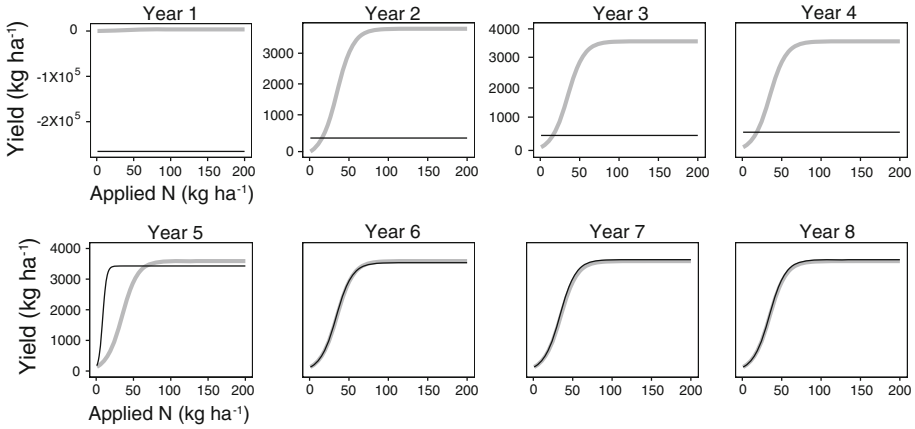

Fig. 9 Comparison of the calculated aggregate yields over the entire field for a range of N values using the true parameter values (*gray lines*), and the parameter posterior distribution means (*black lines*)

low net return realizations due to the lack of parameter convergence (Fig. 10, fourth row and bottom panel). In a real-world scenario it may be advisable for a producer to maintain uniform fertilization except in experimental strips, until a sufficient range of precipitation conditions have been observed (Fig. 10, second panel from the bottom) and some nominal level of convergence is achieved. However, after year 5, the optimizations successfully maximized net returns within the constraint of unknown random precipitation, generating a \$23–25 ha^{-1} advantage over uniform fertilization (Fig. 10, bottom panel). Over time (up to convergence at ~6–10 years) the VRA prescriptions and treatments continually improved, refining knowledge of the driving parameters and increasing profitable returns.

Once the parameters converged (~6 years), the optimization map was largely spatially and temporally static. Different levels of yield were observed, but those levels were driven by the temporal variation of precipitation in our dryland system. If additional information was available on the amount of expected precipitation just prior to fertilization, then the optimization could incorporate that updated information and choose different levels of N. Advance predictions of available water from statistical or process-based models would further increase the accuracy of the optimization and would reduce the risk of over-fertilization especially for such dryland agricultural systems.

Additional variables for future inclusion

The model used for this simulation employs only a few variables for simplicity, however additional driving variables could foreseeably be incorporated into the non-linear function. Whenever explanatory variables are available that help describe the variation in crop yield and reduce spatiotemporal autocorrelation, their use would increase accuracy and further understanding of the system driving factors. For example, if soil water availability was known prior to N fertilizer application, it would likely be more predictive than precipitation in dryland systems with arid moisture regimes. Equally important, if some measure of soil N was available prior to fertilization, then the effects of applied versus intrinsic N could be disentangled, resulting in increased model accuracy. In addition, ecological factors such as weed density and an indicator variable for the previous crop could be included, which would help identify the degree to which the yield-nitrogen relationship depends on plant competition or previous crop nutrient use.

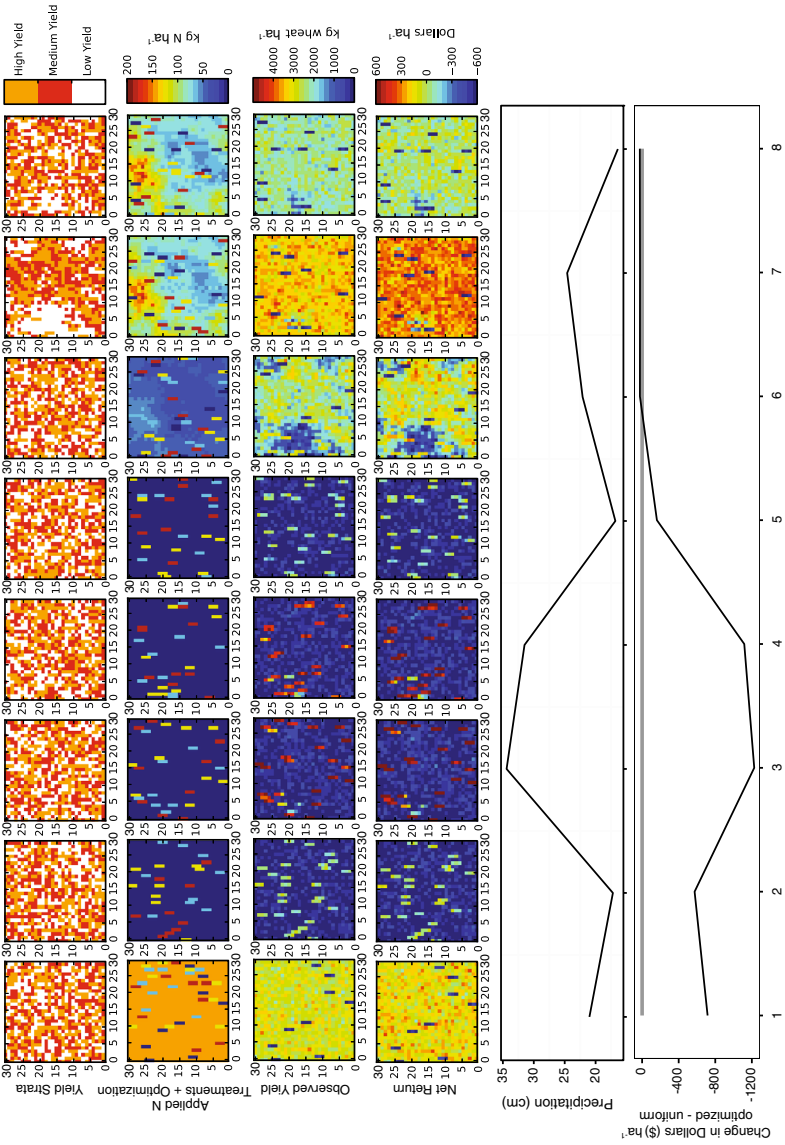

Fig. 10 Simulation of recommended applied N with successive annual model refinement on a theoretical 900 ha field where each 1 ha cell received a prescribed N rate. Details of the sequential process, displayed in columns from left to right, is provided in Table 1. Annual precipitation is included for reference, in addition to the change in dollars/hectare between the optimized and uniform fertilization (140 kg ha^{-1}). Colored blocks in row two represent experimental treatment areas used to reduce time to parameter optimization. Fully optimized maps displayed in the second row, under years 7 and 8

The drawback to including more variables is a possible increase in the number of years of data that would be needed to achieve convergence. Furthermore, with additional variables, more creative parameter space visualizations would be required such as using colors for a fourth dimension. However, the advantage of including key driving variables may offset any drawbacks of additional data requirements, especially if the variables explained large variations in yield.

Conclusion

Our goal was to create an agricultural adaptive management framework that employs VRA and spatial monitoring technologies and explicitly includes the uncertainties of principal factors driving crop yield and quality. While the resulting framework for understanding agricultural management-response relationships has yet to be tested on field-collected data, our results indicate that it has strong potential as a decision aid (with the statistical details hidden) for farmers to progressively manage their nutrient inputs toward an optimization based on maximized net returns. In such a scenario, farmers would generate the data to parameterize the models (experimental designs automatically implemented with VRA), but would then have the ability to modify simulated variables such as precipitation to predict impacts on profitability. Most importantly, this methodology does not require farmers to sacrifice their entire fields to experimentation, as they could simply apply experimental strips until sufficient time had passed to render the data necessary for model convergence.

Another key aspect is that the model incorporates climatic variation, without which prescriptions would be inaccurate, especially in dry-land farming systems. By utilizing a more biologically-appropriate non-linear function instead of the typical linear or linear-quadratic response function (Anselin et al. 2004), the model also moves one step closer towards generating field-based ecological understanding by placing increased emphasis on the local variance rather than the mean crop response to inputs. Furthermore, it simultaneously accounts for spatiotemporal dependence, which will always be a source of bias for regular regression models when farmers are unable to implement the carefully designed experiments commonly used in scientific settings.

In addition to providing an improved model, the entire framework benefits from the inclusion of historical precipitation, input cost and price uncertainties. When farmers make management decisions, they must incorporate all of these uncertainties into their choices if they impact profitability. Therefore, it is only appropriate to quantitatively include uncertainty in the optimizations to provide probabilistic projections of net returns.

In summary, this approach advances precision agriculture towards a probability-based, self-updating system that is consistent with the needs of farmers. It provides adaptability and encourages continual experimentation, leading to an end result of increased profitability, resilience and site-specific understanding of the agricultural system.

Acknowledgments This material is based on work supported by the Montana Institute on Ecosystems' award from the National Science Foundation EPSCoR Track-1 Program under Grant #EPS-1101342 (INSTEP 3). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. This work is also supported by the Montana Fertilizer Advisory Committee and graduate student Grant GW12-004 from the Western Sustainable Agriculture Research and Education Program.

References

Anselin, L., Bongiovanni, R., & Lowenberg-DeBoer, J. (2004). A spatial econometric approach to the economics of site-specific nitrogen management in corn production. *American Journal of Agricultural Economics, 86*(3), 675–687.

Archontoulis, S. V., & Miguez, F. E. (2013). Nonlinear regression models and applications in agricultural research. *Agronomy Journal*. doi:10.2134/agronj2012.0506 .

Baxter, S. J., Oliver, M. A., & Gaunt, J. (2003). A geostatistical analysis of the spatial variation of soil mineral nitrogen and potentially available nitrogen within an arable field. *Precision Agriculture, 4*(2), 213–226.

Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*. 36(2), 192-236.

Besag, J., York, J., & Mollie, A. (1991). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics, 43*(1), 1–59.

Biermacher, J. T., Brorsen, B. W., Epplin, F. M., Solie, J. B., & Raun, W. R. (2009). The economic potential of precision nitrogen application with wheat based on plant sensing. *Agricultural Economics, 40*(4), 397–407. doi:10.1111/j.1574-0862.2009.00387.x.

Bongiovanni, R. G., Robledo, C. W., & Lambert, D. M. (2007). Economics of site-specific nitrogen management for protein content in wheat. *Computers and Electronics in Agriculture, 58*(1), 13–24. doi:10.1016/j.compag.2007.01.018.

Brevik, E. C., Fenton, T. E., & Lazari, A. (2006). Soil electrical conductivity as a function of soil water content and implications for soil mapping. *Precision Agriculture, 7*(6), 393–404. doi:10.1007/s11119-006-9021-x.

Cambardella, C. A., & Karlen, D. L. (1999). Spatial analysis of soil fertility parameters. *Precision Agriculture, 1*(1), 5–14.

Corwin, D. L., & Lesch, S. M. (2003). Application of soil electrical conductivity to precision agriculture. *Agronomy Journal, 95*(3), 455–471.

Corwin, D. L., & Lesch, S. M. (2005). Apparent soil electrical conductivity measurements in agriculture. *Computers and Electronics in Agriculture, 46*(1–3), 11–43. doi:10.1016/j.compag.2004.10.005.

Cressie, N., & Wikle, C. K. (2011). *Statistics for Spatio-Temporal Data*. Hoboken: John Wiley & Sons Inc.

Florin, M. J., McBratney, A. B., & Whelan, B. M. (2009). Quantification and comparison of wheat yield variation across space and time. *European Journal of Agronomy, 30*(3), 212–219. doi:10.1016/j.eja.2008.10.003.

Florin, M. J., McBratney, A. B., Whelan, B. M., & Minasny, B. (2010). Inverse meta-modelling to estimate soil available water capacity at high spatial resolution across a farm. *Precision Agriculture, 12*(3), 421–438. doi:10.1007/s11119-010-9184-3.

Fonnesbeck, C., Patil, A., Huard, D., & Salvatier, J. (2012). *PyMC*. Retrieved June 6, 2013, from http://github.com/pymc-devs/pymc.

Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). *Bayesian Data Analysis* (2nd ed). New York, NY: Chapman & Hall/CRC.

Jackson, G. (1998). Predicting Spring Wheat Yield and Protein Response to Nitrogen. *MSU Fertilizer Facts,* no. 17. Retrieved August 10, 2013 from http://www.sarc.montana.edu/php/Research/ffacts/?id=17

Jiang, P., He, Z., Kitchen, N. R., & Sudduth, K. A. (2009). Bayesian analysis of within-field variability of corn yield using a spatial hierarchical model. *Precision Agriculture, 10*(2), 111–127. doi:10.1007/s11119-008-9070-4.

Jung, W. K., Kitchen, N. R., Sudduth, K. A., Kremer, R. J., & Motavalli, P. P. (2005). Relationship of apparent soil electrical conductivity to claypan soil properties. *Soil Science Society of America Journal, 69*(3), 883–892.

Kerry, R., & Oliver, M. A. (2003). Variograms of ancillary data to aid sampling for soil surveys. *Precision Agriculture, 4*(3), 261–278.

Khosla, R., Inman, D., Westfall, D. G., Reich, R. M., Frasier, M., Mzuku, M., Koch.B. and Hornung.A. (2008). A synthesis of multi-disciplinary research in precision agriculture: site-specific management zones in the semi-arid western Great Plains of the USA. *Precision Agriculture, 9*(1–2), 85–100. doi:10.1007/s11119-008-9057-1.

King, J. A., Dampney, P. M. R., Lark, R. M., Wheeler, H. C., Bradley, R. I., & Mayr, T. R. (2005). Mapping potential crop management zones within fields: use of yield-map series and patterns of soil physical properties identified by electromagnetic induction sensing. *Precision Agriculture, 6*(2), 167–181.

Koch, B., Khosla, R., Frasier, W. M., Westfall, D. G., & Inman, D. (2004). Economic feasibility of variable-rate nitrogen application utilizing site-specific management zones. *Agronomy Journal, 96*(6), 1572–1580.

Kravchenko, A. N., Robertson, G. P., Thelen, K. D., & Harwood, R. R. (2005). Management, topographical, and weather effects on spatial variability of crop grain yields. *Agronomy Journal, 97*(2), 514–523.

Kühn, J., Brenning, A., Wehrhan, M., Koszinski, S., & Sommer, M. (2008). Interpretation of electrical conductivity patterns by soil properties and geological maps for precision agriculture. *Precision Agriculture, 10*(6), 490–507. doi:10.1007/s11119-008-9103-z.

Laird, N. M., & Ware, J. H. (1982). Random effects models for longitudinal data. *Biometrics, 38*, 963–974.

Lambert, D. M., Lowenberg-DeBoer, J., & Malzer, G. L. (2006). Economic analysis of spatial-temporal patterns in corn and soybean response to nitrogen and phosphorus. *Agronomy Journal, 98*(1), 43. doi:10.2134/agronj2005.0005.

Lichstein, J. W., Simons, T. R., Shriner, S. A., & Franzreb, K. E. (2002). Spatial autocorrelation and autoregressive models in ecology. *Ecological Monographs, 72*(3), 445–463.

Liu, Y., Swinton, S. M., & Miller, N. R. (2006). Is site-specific yield response consistent over time? Does it pay? *American Journal of Agricultural Economics, 88*(2), 471–483.

Mamo, M., Malzer, G. L., Mulla, D. J., Huggins, D. R., & Strock, J. (2003). Spatial and temporal variation in economically optimum nitrogen rate for corn. *Agronomy Journal, 95*(4), 958–964.

Meyer-Aurich, A., Weersink, A., Gandorfer, M., & Wagner, P. (2010). Optimal site-specific fertilization and harvesting strategies with respect to crop yield and quality response to nitrogen. *Agricultural Systems, 103*(7), 478–485. doi:10.1016/j.agsy.2010.05.001.

Montana Wheat and Barley Committee. (2013). Pricing: Montana Wheat & Barley Committee. Retrieved April 4, 2013, from http://wbc.agr.mt.gov/wbc/Producers/Pricing/local/2013_GreatFalls.xls.

National Climatic Data Center. (2013). NCDC: Precipitation Data. Retrieved March 3, 2013, from http://gis.ncdc.noaa.gov/map/viewer/#app=cdo&cfg=cdo&theme=precip&layers=000111.

Patzold, S., Mertens, F. M., Bornemann, L., Koleczek, B., Franke, J., Feilhauer, H., et al. (2008). Soil heterogeneity at the field scale: a challenge for precision crop protection. *Precision Agriculture, 9*(6), 367–390. doi:10.1007/s11119-008-9077-x.

R Core Team. (2012). *R*: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from http://www.R-project.org/.

Robinson, G. K. (1991). That BLUP is a good thing: the estimation of random effects. *Statistical Science, 6*(1), 15–32.

Sadler, E. J., Sudduth, K. A., & Jones, J. W. (2007). Separating spatial and temporal sources of variation for model testing in precision agriculture. *Precision Agriculture, 8*(6), 297–310. doi:10.1007/s11119-007-9046-9.

Schlather, M. (2012). RandomFields: Simulation and Analysis of Random Fields. Retrieved November 6, 2012, from http://CRAN.R-project.org/package=RandomFields.

Shahandeh, H., Wright, A. L., & Hons, F. M. (2010). Use of soil nitrogen parameters and texture for spatially-variable nitrogen fertilization. *Precision Agriculture, 12*(1), 146–163. doi:10.1007/s11119-010-9163-8.

Shahandeh, H., Wright, A. L., Hons, F. M., & Lascano, R. J. (2005). Spatial and temporal variation of soil nitrogen parameters related to soil texture and corn yield. *Agronomy Journal, 97*(3), 772. doi:10.2134/agronj2004.0287.

Thöle, H., Richter, C., & Ehlert, D. (2013). Strategy of statistical model selection for precision farming on-farm experiments. *Precision Agriculture*. doi:10.1007/s11119-013-9306-9.

Thrikawala, S., Weersink, A., Fox, G., & Kachanoski, G. (1999). Economic feasibility of variable-rate technology for nitrogen on corn. *American Journal of Agricultural Economics, 81*(4), 914–927.

USDA ERS. (2012). Commodity Costs and Returns. Retrieved April 7, 2013, from http://www.ers.usda.gov/datafiles/Commodity_Costs_and_Returns/Data/Current_Costs_and_Returns_All_commodities/cwhea.xls.

USDA ERS. (2012). Fertilizer Price Indexes, 1960-2012. Retrieved April 7, 2013, from http://www.ers.usda.gov/datafiles/Fertilizer_Use_and_Price/Fertilizer_Prices/table7.xls.