OPTIMIZING SITE-SPECIFIC ADAPTIVE MANAGEMENT USING A PROBABILISTIC FRAMEWORK: EVALUATING MODEL PERFORMANCE USING HISTORIC DATA

P.G. Lawrence, L.J. Rew, B.D. Maxwell

Land Resources and Environmental Sciences Dept.
Montana State University
Bozeman, Montana

ABSTRACT

Agricultural producers are tasked with managing crop yield responses to nitrogen (N) within systems that have high levels of spatial (biophysical), climatic, and price uncertainty. To date, the outcome of most variable rate application (VRA) research has focused on the spatial dimension, proposing optimal fertilizer prescription maps that can be applied year after year. However, temporally static prescriptions can result in suboptimal outcomes, particularly if they do not consider the impact and likelihood of alternative weather or price regimes that can drastically alter crop responses and net returns. Furthermore, most optimizations are built on the assumption of linear crop responses when non-linearity may be more biologically appropriate and could result in altered N prescriptions.

In this presentation, we outline our methodology to address these uncertainties using a non-linear spatiotemporal Bayesian updating framework. This strategy continually improves N optimizations, increases net returns and reduces uncertainty in the parameter estimates. The framework is able to quantify the probabilities of different net return outcomes, allowing the producer to choose their N management based on their particular level of risk adversity. It also enables the producer or researcher to assess the impacts of future scenarios such as prolonged drought or price fluctuations.

This methodology was tested within a simulation to assess the number of years required for model convergence and enhanced net returns. It was then applied to the years 1980 – 1992 to hindcast the impact of extended drought in Montana during 1987-1991. Simulated crop responses incorporated realistic levels of residual variability based on ten years of observations from a dryland wheat farm located near Great Falls, Montana. For simplicity, the crop was assumed to respond non-linearly to variation in soil apparent electrical conductivity (EC_a), applied nitrogen (N), and precipitation. Historical wheat price data from this region also informed the model and served as an additional source of variability that impacted the net returns.

Parameter convergence and net returns higher than those of uniform fertilization were achieved after six to eight years, resulting in a spatial net return benefit of $23-25/hectare. After year six, the spatial random effects in the model effectively eliminated the confounding influence of spatial autocorrelation on the crop response coefficients. Small experimental N rate treatments (0, 60, 120, 180 kg/ha) were randomly applied each year as a part of this framework to ensure that crop responses to N were explored under the full space of possible soil and precipitation conditions. These strip experiments reduced the time required for convergence of the parameter estimates.

During the late 1980s, the severe drought in Montana reduced hypothetical savings from a level of $450,000 in 1983 to below zero as early as 1988. The impacts on savings are mirrored in governmental data on farm bankruptcies during this period. Substantial variability remained around the estimates for the different fertilization scenarios; however the optimized fertilizer prescriptions consistently out-performed the uniform prescriptions on a field-wide basis. With a nominal level of governmental price support, producers spatially optimizing their N inputs would have survived the drought. Producers applying uniform levels of fertilizer would have increased levels of debt, especially under low and high input levels.

This simulation study demonstrated a useful decision aid framework that can empower agricultural producers with site-specific management that accounts for the range of possible uncertainties producers must face. Decision support tools must be applicable across years rather than being optimal under only one set of climatic conditions. Decision support tools must use crop response functions that are biologically appropriate yet statistically tractable. Finally, the decision aid must acknowledge the variability not only in crop responses, but also the variability in crop prices that has a strong impact on net returns and management strategies. With the uncertainty associated with future climates, an approach for monitoring system agronomic and economic performance is crucial for maintaining resilient agro ecosystems. The framework developed here meets all of these requirements and can be easily adapted to incorporate additional driving variables or alternative crop response functions. By providing a flexible platform for progressively refining system parameters and optimizing spatial N prescriptions, this research provides a baseline tool that may be useful to producers across a wide range of crops and growing conditions.

Keywords: variable rate application; Bayesian statistics; input optimization; simulation experiment; decision support; spatial variation

INTRODUCTION

Dryland farmers throughout the world are forced to manage spatial, bioclimatic, and economic variability on multiple scales. Despite these diverse management challenges, most agricultural research in dryland regions is focused on optimizing individual agronomic decisions while minimizing the confounding effects of extraneous factors (Suppe, 1987; Cook *et al.*, 2013). While isolating best management practices is a worthwhile endeavor, such an approach makes no effort to reconcile the numerous on-farm uncertainties that ultimately determine the economic and environmental sustainability of specific farms.

In order to contextualize such research and ensure its practical relevance, it must be integrated with the other forms of variability at scales larger than the field. To do so requires a framework that can synthesize multiple data streams and produce a probabilistic estimation of the impact of individual management strategies. Thus our research question was: Can we construct a model framework that can arrive at optimized variable rate N fertilizer recommendations for dryland spring wheat given different climate and price scenarios?

Such a decision-support and statistical framework must be able to:

1. Assimilate site-specific climatic, economic and environmental variables into a model that is based on biologically meaningful relationships (i.e. site and history specific).
2. Progressively improve the knowledge of systemically-important parameters over time.
3. Produce management prescriptions for each year that make the best use of past and current information and serves as an experiment within the next growing season.
4. Perform the above while accounting for spatiotemporal variation and spatial autocorrelation.

We propose such a framework that merges precision agriculture data, historical economic records, localized precipitation measurements, and on-farm experimentation within a Bayesian statistical framework. Each growing season, a cycle of experimentation, observation, and synthesis leads to sequentially greater understanding of the agroecosystem's driving parameters and how they influence crop performance on small and large spatial scales. Incorporating multiple years minimizes the confounding influence of climatic variation, which is effectively ignored when only one or two years of data are used (Anselin *et al.*, 2004; Shahandeh *et al.*, 2005; Liu *et al.*, 2006).

Adopting a Bayesian probabilistic approach enables each producer to weigh their level of risk preference against the probabilities of realizing specific weather or economic outcomes when deciding between alternative management options. This is performed within the context of fertilizer management, however extension to other management decisions and sustainability objectives is anticipated.

By using a multi-year decision-support system, hypothetical future scenarios can easily be explored. To assess the impact of climatic variation or climatic change, a farmer could explore the economic implications of prolonged droughts or wet periods. Within this paper, we examine these impacts within the context of a historic drought period during the late 1980s, which demonstrates how the framework could be applied to future climatic scenarios.

METHODS

model construction

The agronomic yield model underlying this framework integrates quantitative, interacting measures of edaphic, nitrogen (N), and precipitation variation. Specifically, the relationship between these variables and yield follows a logistic form (Archontoulis and Miguez, 2013):

$$Yield_{ij} = \frac{\beta max * precip_j}{1+\exp(\beta_{Shp} - \beta_1 * QuantN_{ij} - \beta_2 * EC_{a,i} - \beta_3 * EC_{a,i} * QuantN_{ij})} + \phi_i + \varepsilon \qquad (1)$$

Where $\varepsilon \sim N(0,\sigma_e^2)$, and i is the spatial effect of cell I, EC_a represents electrical conductivity measurements as a proxy for edaphic variation, $QuantN_{ij}$ is the amount of applied N in kg/ha, and *precip* is the water year precipitation in cm. EC_a reflects a suite of soil texture-related conditions, however these soil properties directly influence plant yield, thus it is still an adequate method for characterizing soil variation within this context (Corwin and Lesch*, 2003; Jung et al.*, 2005; King *et al.*, 2005). The parameter β_{max} represents the maximum amount of yield at the asymptote, and β_{shp} represents a shape parameter.

This functional form asserts that the maximum yield possible in any given year is determined by the level of moisture availability. Including N and EC_a within the exponential term of the denominator characterizes the interaction of edaphic variation with management inputs, which are then moderated by the annual level of precipitation.

The output from the yield model was then incorporated into a net return function that integrates the data, model, and parameter uncertainties (**Fig. 1**). Uncertainty in the price of N and in the precipitation is included by selecting random draws from the historical distributions of these variables. The price of wheat experienced by the farmer was included by selecting random draws, 365 days in advance, from the posterior distribution of an autoregressive lag one (AR1) model of the first-differenced historical prices obtained from the Montana Wheat and Barley Committee ("Pricing :: Montana Wheat & Barley Committee" , 2013). Wheat prices were suitable for an AR1 model due to their high temporal density and strong short-term autocorrelation, however N prices were only

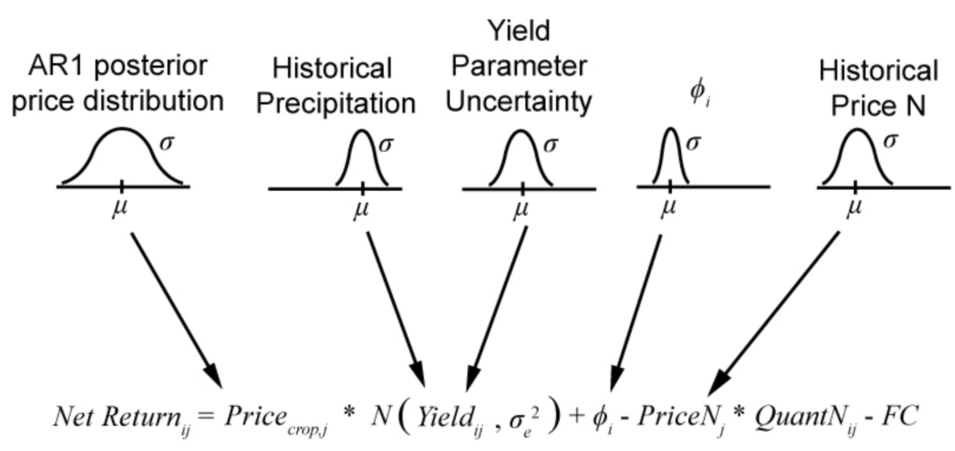

Fig. 1. Schematic for the inclusion of the wheat price, precipitation, yield, N price, model, parameter, and spatial uncertainty into the net return function.

available on a yearly basis ("USDA ERS - Fertilizer Price Indexes, 1960-2012", 2012), as were measures of growing season precipitation (site: Sun River 4s; National Climatic Data Center, 2013). Therefore, the precipitation distribution was approximated by a normal representation of its historical distribution $N(\mu, \sigma)$ = N(26.2 cm/growing season, 6.4 cm/growing season), and the N distribution was similarly represented by historical values $N(\mu, \sigma)$ = N($0.55/kg, 0.055$/kg). *FC* represents the other average fixed costs associated with crop management ($605.44/ha; "USDA ERS - Commodity Costs and Returns," 2014). To account for spatial autocorrelation, the spatial random effect ϕ_i for each cell was included, which specifies a Conditional AutoRegressive (CAR) model (Jiang *et al.*, 2009) as detailed below.

yield model priors

The priors used for the CAR and the non-linear model follow the suggestions of Gelman *et al.* (2004) and Jiang (2009). Parameters within the yield model used diffuse priors of $N(0, 1000)$, and dispersion parameters were modeled with non-informative Inverse-Gamma priors ~ $IG(.01,100)$. In contrast to a typical regression model, the CAR model adds a spatial random effect to the model for the mean, with each cell conditionally dependent on the neighboring cells:

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

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

where ∂_i represents the set of neighbors surrounding cell i, and m_i is the number of these neighbors (Banerjee *et al.*, 2004). This implies that Y_i is determined both by the value of the explanatory variables but also by the values of adjacent yield values.

optimization

To obtain annual N prescription maps, a net return function must be available for optimization. Using a Bayesian approach, the net return function would consist of distributions rather than individual functions, prohibiting such an approach. Therefore, a Monte-Carlo method was chosen for taking random draws from the distributions of input parameters to obtain many different realizations of the net return function. The net return functions were then optimized, forming a distribution of optimal N values for which profit could be maximized. The final optimal (over space and time) N distributions for each cell thus incorporated the entire uncertainty of the agricultural system in order to achieve a recommendation.

experimentation and parameter space exploration

Ensuring parameter convergence in a reasonable amount of time requires that a wide range of independent variable combinations are observed. For example, if only one N rate was applied to a field every year, then there would be no data to support conclusions on the yield response to alternative N rates. Therefore, exploring the N-EC_a-precipitation parameter space as efficiently as possible is critical for parameter convergence and optimization. Unfortunately, only the N variable is subject to manipulation, and the precipitation variable is unpredictable. Nevertheless, N rates can be applied across multiple years in areas with different EC_a values to effectively explore the parameter space in as little time as possible given the precipitation uncertainty. To accomplish this, after each year's yields were observed, the cells in the field were stratified into three different yield classes, within which different N rates were randomly applied. This procedure was automated as a component of the framework.

annual updating

A core advantage of the Bayesian approach is the ability to easily update parameter estimates when new data are received. This advantage is particularly helpful within agricultural systems due to the annual nature of observations and the ability to perform input manipulations each year. Therefore, each year's inputs and observations can be regarded as an experiment that continually updates knowledge about the location and precision of the system's driving variables.

Within this framework, the annual experiment consists of the net return-maximizing prescription map in conjunction with the parameter-space experiments. These spatial input data were then matched to the observed yield data to further refine knowledge of β_{max}, β_1, β_2, and β_3.

model implementation and simulation

To assess the ability of this model to converge and provide useful prescriptions and forecasts, a 30 by 30 cell grid was created, with each cell representing a hectare. Although this grid size may not be realistic from a practical perspective, it is useful for demonstration purposes, and the inferences could be easily scaled down to any desired size.

Observed yields and yield responses were based on data from a non-irrigated wheat (*Triticum aestivum*) agricultural system near Great Falls, Montana. The spatially correlated EC_a Gaussian Random Field grid was generated within the R package **RandomFields** (R Core Team, 2012; Schlather, 2012) and was characterized by an exponential spatial covariance structure (σ^2=640, μ = 50, range=50, nugget=0, scale=1).

Initial conditions for the simulated updating process assumed that a farmer beginning to use PA technology would start with at least one 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 numbers of observations (high, medium, low), within which different N rate treatments were applied. Choosing an equal number of observations ensured that each class, representing cells with different productivity potentials, would contain sufficient

Table 1. "Real" parameters used to calculate yield within equation 6. 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 45000 (parameterized as an inverse in python package pymc (Fonnesbeck *et al.*, 2012) as .000022 (1/45000)).

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

data to characterize the unique yield response and to implement three repetitions of the N rate experiment. 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. This fertilization experimentation system was automated and has been used on real farm fields.

To calculate yields in the initial year and in subsequent iterations, *equation (1)* was applied using the parameter coefficients (**Table 1**). 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). 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 yield experiment. 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 *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 (5625 kg/ha σ^2) for neighboring cells, and 0 kg/ha for non-neighboring cells. These values were based on observed spatial correlations from the previously mentioned field experiment.

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.

Priors used for the coefficient parameters followed either normal or truncated normal distributions (Table 2; 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

Table 2. 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/σ^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/σ^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/σ^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/σ^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/σ^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_τ)	*Gamma*(0.01, 0.01)

that could 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 for the total model variance, σ_e^2 and the spatial variance (τ_c^2) were set to follow inverse-gamma distributions (~$IG(a,b)$; Jiang *et al.*, 2009), which were again specified to be non-informative (Gelman *et al.*, 2004). During each year, the model was run for 100,000 iterations, using a burn-in period of 40,000 samples and a thinning rate of 20 in order to improve convergence and reduce autocorrelation between the samples.

model hindcasting assessment

To assess the ability of the model to provide useful estimates of economic resilience, the MCMC-derived yield model was applied to the same field under different N strategies, using the same constant EC_a values. Instead of simulating precipitation, values were gathered from historical growing season accumulations during the years 1980-1992 in the same location near Great Falls, MT. Wheat and N prices were obtained for the same period from the USDA ("USDA ERS - Wheat: Farm Prices, Support Prices, and Ending Stocks, 2006; "USDA ERS - Fertilizer Price Indexes, 1960-2012", 2013). The initial savings of the farmer were assumed to be \$100,000, after which the different N inputs were applied each year. To adequately capture the spread in yield and therefore economic outcomes that would result from the variability in the yield model parameters, 100 repetitions were run for all cells, in each year, and under each N input scenario. The results were characterized by first averaging and calculating the standard deviations for the total yield in each cell. Subsequently, these average and standard deviation values were summed for the entire field and then piped into the net return function to obtain ranges of economic outcomes.

RESULTS AND DISCUSSION

Model diagnostic plots (**Fig 2**), indicate that parameter convergence was achieved for β_{max}, β_1, β_2 and β_3 after six years, with convergence achieved for σ_ε after eight years and σ_s (spatial variance parameter) approaching convergence after eight years. Repeat simulations indicate that this time to convergence is consistent. The long time to convergence suggests that creating optimal prescription maps only using several years of data will be misleading due to the temporally short sampled span of climate data. Therefore, ensuring that the dataset utilizes a range of observed precipitation or climatic conditions is vitally important to accurately understanding the conditional crop responses.

The parameter-space plots (e.g. **Fig 3**) indicate that the N-rate experiments were successful at exploring the N-EC_a-precipitation parameter space. Had only a uniform level of N or an optimized N rate been applied, the observations would have been clustered in one region of the parameter space, increasing the time to convergence or even preventing it.

Residuals extracted from year six of the simulation display no spatial pattern, suggesting that spatial autocorrelation was sufficiently managed within

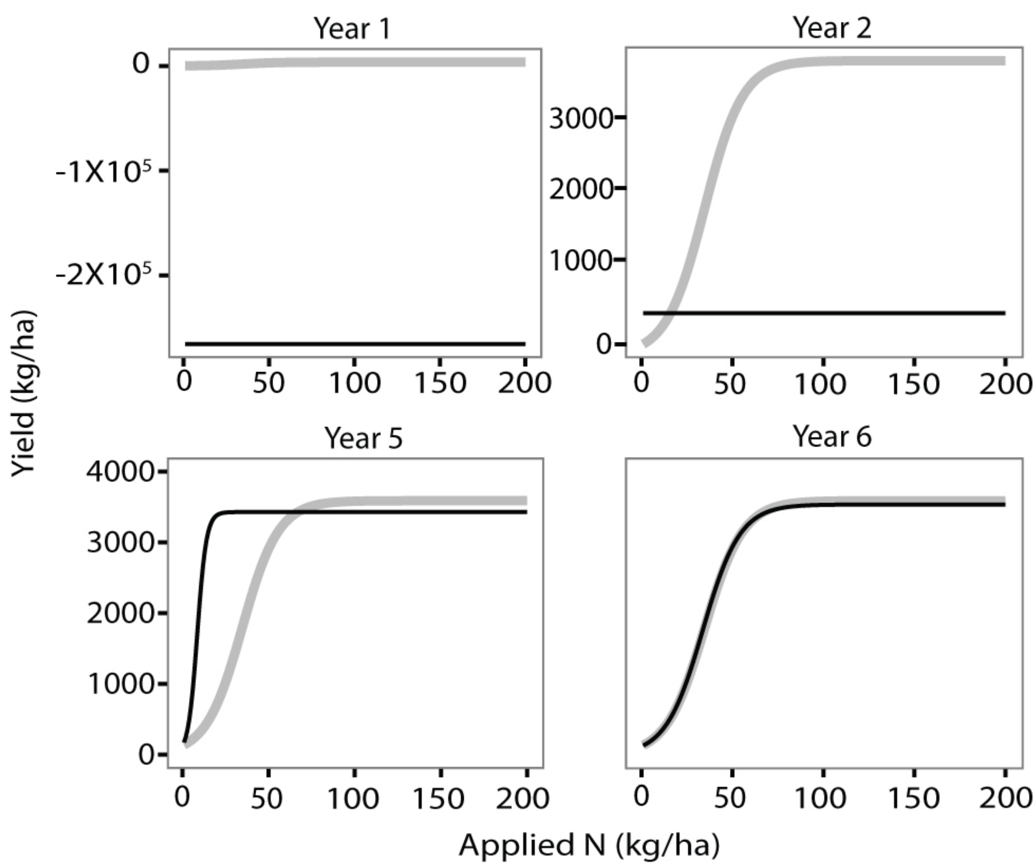

Fig. 2. Convergence of the yield values estimated using the regression model parameters (black) versus the yield values generated from the true, known parameters (gray).

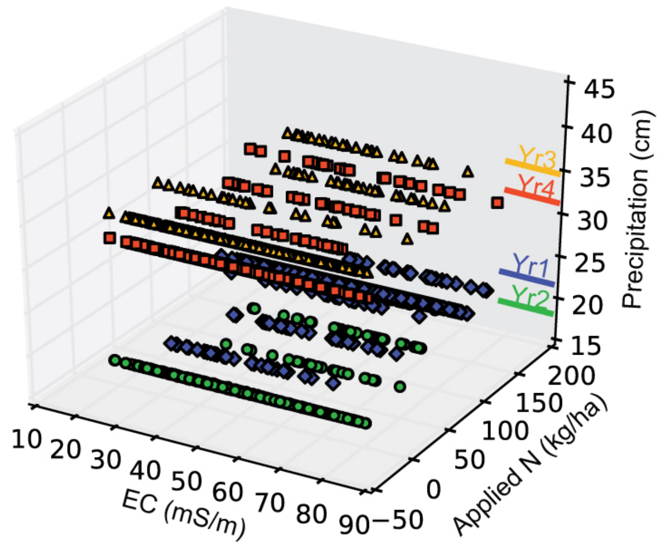

Fig. 3. Plot of the realized N-EC-precipitation values for all cells in years one through four.

the model (Moran's I = 0.01, p-value for significant spatial autocorrelation=0.48).

A visual assessment of the stratification, updating, and optimization process (**Fig 4**) shows that the optimal N levels stabilized after six years of data were collected. Before year six, the optimized N treatments performed far worse (not shown) simply because the parameter estimates had not converged, thus a rate of 0 kg/ha was selected by the optimizing function, resulting in net losses. From a practical standpoint, before the sixth year it would be advised for a farmer to maintain uniform levels outside of the experimental areas in order to retain profitability. After convergence, the net returns from the optimized N inputs were $23-25 dollars/ha higher than net returns resulting from uniform management at 120 kg/ha. In any given year, some of the cells receiving the uniform treatment would likely perform better, however on a whole field basis the optimal N strategy outperformed the uniform strategy across climatic scenarios.

The predictions of the framework regarding economic responses to drought indicated that the years 1987-1992 were indeed disastrous for farmers in the Northern Great Plains. None of the management strategies were able to retain savings above the level of zero dollars (**Fig 5**), and it is well documented that many producers during this time period either were forced into bankruptcy or received substantial governmental assistance. Chapter 12 bankruptcy filings for the mountain states increased from 3.42 to 44.79 per 10,000 farms between 1986 and 1987, then stayed relatively high through 1992 (**Fig 6**; Stam and Dixon, 2004). Some of this spike may be explained by introduction of the Chapter 12 bankruptcy law in 1986, which specifically favored farmers seeking to file for bankruptcy. However, the magnitude and coincidence with the drought period suggests that the lack of precipitation may nevertheless have had a strong effect.

Despite the impact of drought, farmers who adopted a spatially optimized N management strategy (albeit not technologically feasible during the 1980s)

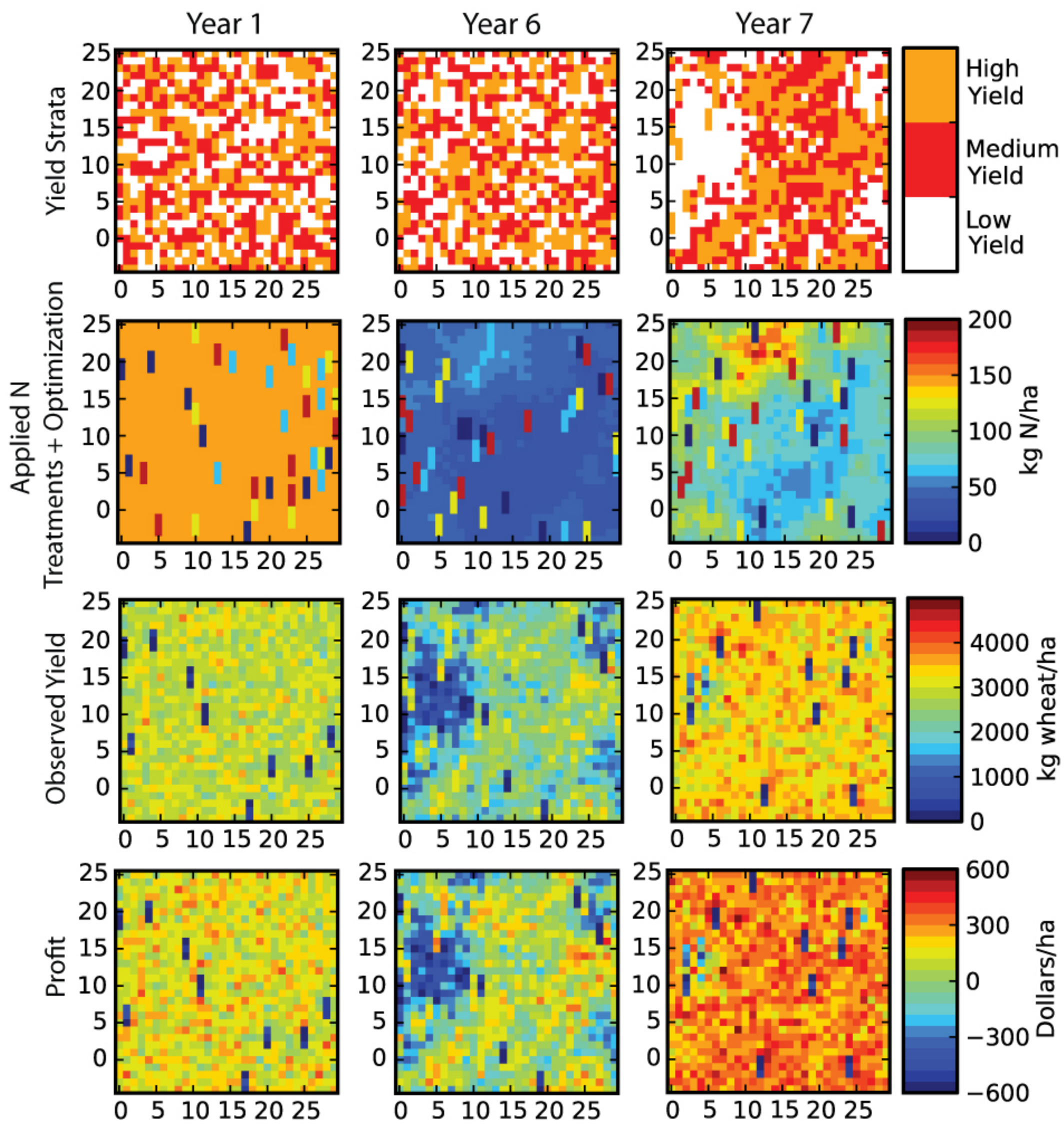

Fig. 4. Demonstration of the stratification process, experimental layout, yield and profit calculation for years 1, 6, and 7. Stable optimization was achieved after the data for year six was collected, and remained stable thereafter.

would have outperformed all other management strategies. The uncertainty of these predicted outcomes increased as the time horizon of the forecast became longer, but the average savings trends nevertheless captured the consistent disparity between the fertilization strategies.

Applying the predictive strategy used here to forecasted precipitation or economic trends instead of using historical data could have great utility for planning agricultural adaptation or creating policy to mitigate climate change. Similar adaptations could also be planned for price fluctuations or other disturbances by developing novel scenarios from historical data. Most importantly, these forecasted outcomes are farm-specific, thus they can be easily modified for every location and will demonstrate a large range of profitability

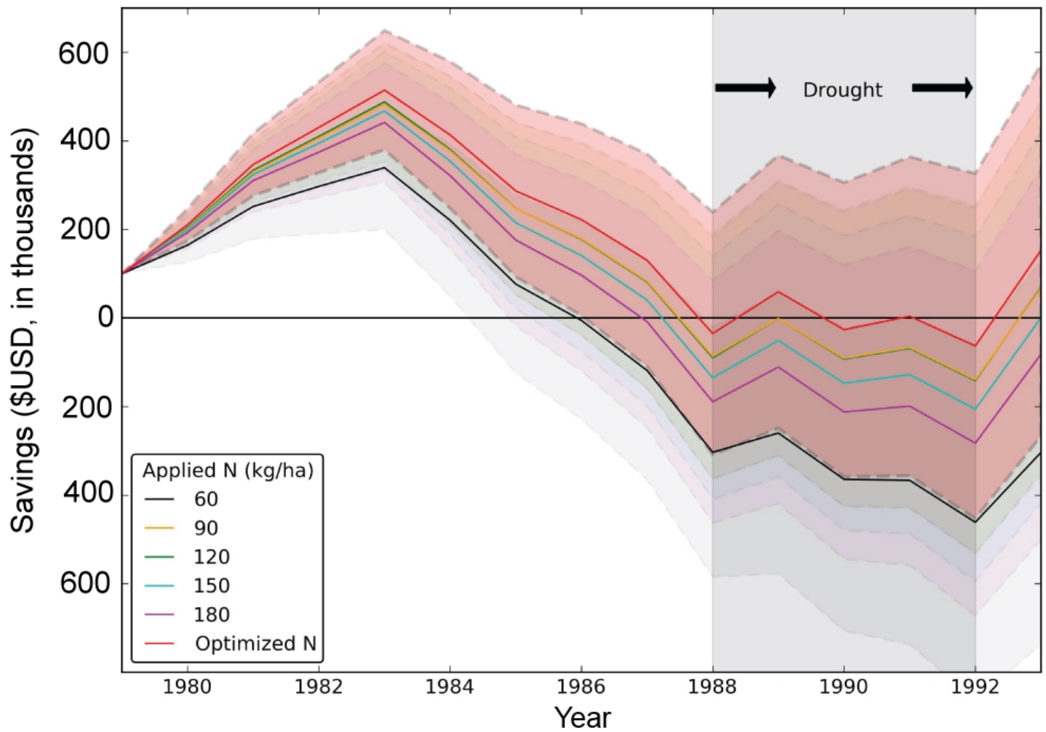

Fig 5. Simulated net savings trajectory under different N input levels during 1980-1992. Shaded areas represent one standard deviation confidence bands for assumed variability in yield responses.

outcomes for different bioclimatic regions.

Thus, our decision support system could be adapted for assessing all purchased inputs on a field by field basis, overcoming much of the uncertainty in applying

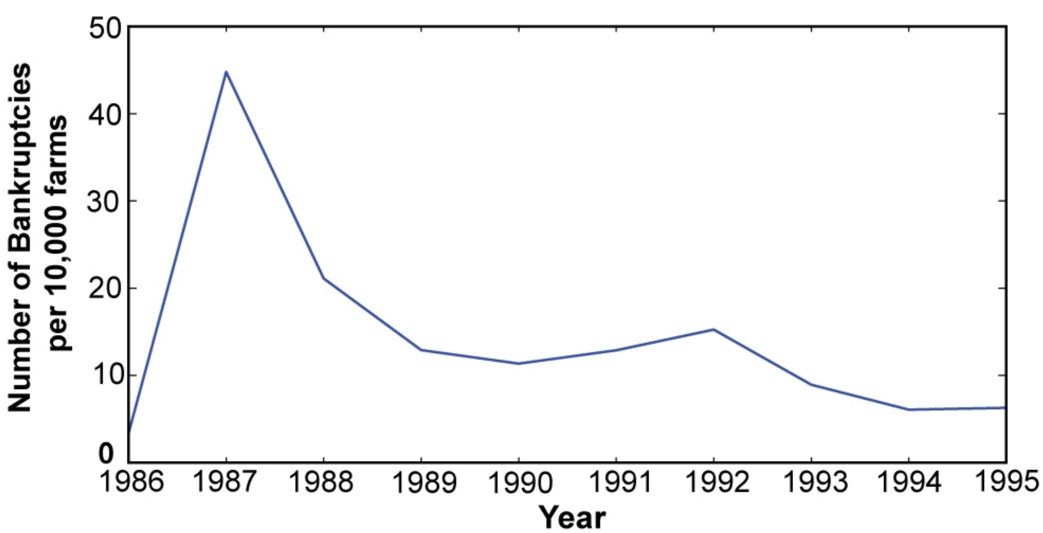

Fig 6. Number of chapter 12 farm bankruptcies per 10,000 farms in the region containing mountain states AZ, CO, ID, MT, NV, NM, UT and WY. Data from Stam and Dixon (2004).

standard research center-based agricultural research.

CONCLUSION

The ability of PA technologies to collect large quantities of spatiotemporal data continues to progress rapidly. With this profusion of data sources, significant effort will be devoted to optimizing individual components of farming systems in isolation from other confounding elements. What is also needed is an ability to contextualize each of these components in order to understand their relative importance to the farm as a whole. While large-scale farming becomes more technologically driven and compartmentalized, simultaneous efforts are necessary to give farmers an understanding of how their management choices influence their farm's overall economic and ecological resilience.

The framework presented here provides a first step towards probabilistically integrating site-specific management, soils, climatic, and economic data in a tool for optimizing N management and predicting future economic outcomes. Results from the model diagnostics, optimization plots, and scenario testing reinforce its potential as a tool that can be applied to a variety of farms and bioclimatic scenarios. By incorporating a non-linear yield model, accounting for spatiotemporal correlation, and adopting a Bayesian approach, the framework offers a significant improvement over previous economic optimization methods. Nevertheless, opportunities for improvement remain, especially for estimating and balancing environmental with economic outcomes, such as minimizing N losses to the soil.

Models are inherently simplifications of reality. In constructing this framework, we have attempted to choose the appropriate number of driving variables that are required for understanding the dynamics of agricultural systems. Future extensions of the framework will likely include sub-components for estimating ecological externalities or economic variables such as crop insurance payments. Without these extensions, the current scaffolding for integrating these disparate components nevertheless enables flexibility while providing the probabilistic grounding for understanding the future impacts of perturbations, stresses and management choices. If integrated into a farmer-oriented web application, such a framework could be a valuable adaptive decision-making tool.

Climate change and economic variability are two of the most significant disturbances that will likely impact agricultural systems in the next 50 years. As demonstrated with data from the 1980s, this framework has the ability to help predict future economic outcomes, thus it has strong potential for probabilistically estimating the impacts of these more severe disturbances. Only then will it be possible to understand how current management decisions impact the future resilience of agricultural systems, and how those decisions may be altered to reach a more sustainable future.

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**: 675–687.

Archontoulis S.V. & Miguez F.E. 2013. Nonlinear Regression Models and Applications in Agricultural Research. *Agronomy Journal*, **105**: 1-13.

Banerjee S., Carlin B.P., and Gelfand A.E. 2004. *Hierarchical Modeling and Analysis for Spatial Data.* Chapman & Hall/CRC, New York, NY.

Cook S., Cock J., Oberthur T., and Fisher M. 2013. On-Farm Experimentation. *IPNI Better Crops*, **97**: 17–20.

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

Fonnesbeck C., Patil A., Huard D., and Salvatier J. 2012. *PyMC* [software]. Available from https://github.com/pymc-devs/pymc.

Gelman A., Carlin J.B., Stern H.S., and Rubin D.B. 2004. *Bayesian Data Analysis.* Chapman & Hall/CRC, New York, NY.

Jiang P., He Z., Kitchen N.R., and Sudduth K.A. 2009. Bayesian analysis of within-field variability of corn yield using a spatial hierarchical model. *Precision Agriculture*, **10**: 111–127.

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

King J.A., Dampney P.M.R., Lark R.M., Wheeler H.C., Bradley R.I., and 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**: 167–181.

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

Pricing :: Montana Wheat & Barley Committee. 2013. Retrieved June 15, 2013, from http://wbc.agr.mt.gov/wbc/Producers/Pricing.html

R Core Team. 2012. *R: A Language and Environment for Statistical Computing* [software]. R Foundation for Statistical Computing, Vienna, Austria. Available from http://www.r-project.org.

Schlather M. 2012. *RandomFields: Simulation and Analysis of Random Fields* [software]. Available from http://cran.r-project.org/web/packages/RandomFields/index.html.

Shahandeh H., Wright A.L., Hons F.M., and Lascano R.J. 2005. Spatial and Temporal Variation of Soil Nitrogen Parameters Related to Soil Texture and Corn Yield. *Agronomy Journal*, **97**: 772.

Suppe F. 1987. The limited applicability of agricultural research. *Agriculture and Human Values*, **4**: 4–14.

Western Regional Climate Center. 2014. Monthly Precipitation Data for Station Sun River 4S, 1979-1992. Retrieved March 20, 2014, from http://www.wrcc.dri.edu/cgi-bin/cliMAIN.pl?mt8021.

USDA ERS. 2006. Wheat: Farm Prices, Support Prices, and Ending Stocks. Retrieved April 2, 2013, from http://ers.usda.gov/datafiles/Wheat_Wheat_Data/Static_Snapshots/2006wheatyrbkapptable9_1_.xls

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

USDA ERS. 2014. Historical Wheat Commodity Costs and Returns. Retrieved March 10, 2014, from http://ers.usda.gov/datafiles/Commodity_Costs_and_returns/Data/Historical_costs_and_returns_Wheat/HUSWhea.xls