The International Society of Precision Agriculture presents the
15th International Conference on Precision Agriculture
26–29 JUNE 2022
Minneapolis Marriott City Center | Minneapolis, Minnesota USA

OPTIMIZING NITROGEN APPLICATION TO MAXIMIZE YIELD AND REDUCE ENVIRONMENTAL IMPACT IN WINTER WHEAT PRODUCTION

A. Peerlinck[1], G. Morales[1], J. Sheppard[1], P. Hegedus[2], B. Maxwell[2]

[1]Gianforte School of Computing, Montana State University, Bozeman, MT.

[2]Land Resources & Environmental Science, Montana State University, Bozeman, MT.

A paper from the Proceedings of the
15th International Conference on Precision Agriculture
June 26-29, 2022
Minneapolis, Minnesota, United States

Abstract.

Field-specific fertilizer rate optimization is known to be beneficial for improving farming profit, and profits can be further improved by dividing the field into smaller plots and applying site-specific rates across the field. Finding optimal rates for these plots is often based on data gathered from the plots, which are used to determine a yield response curve, telling us how much fertilizer needs to be applied to maximize yield. In this research, we trained a Random Forest to create plot-specific non-parametric yield response curves. We then use these curves to determine the optimal amount of fertilizer to be applied to specific plots by maximizing a net return function based on these curves. However, we claim that there are additional issues that should be taken into account when designing optimal prescription maps. In addition to optimizing yield, we want to reduce strain on farming equipment by minimizing rate jumps between consecutive cells. This helps machines run more efficiently and last longer, thus reducing waste. Furthermore, when creating these optimized prescription maps, we also aim to improve environmental impact by reducing the overall fertilizer applied, as excess nitrogen seeps into the soil and drains into our waterways, negatively affecting water quality. In previous work, we found that it is possible to reduce overall fertilizer applied by 5 to 10% when creating experimental prescription maps without significantly impacting yield. Therefore, we hypothesize this will hold true for optimized prescription maps as well. We address these three separate, competing objectives using an adjusted genetic algorithm, known as Non-Dominated Sorting Genetic Algorithm II (NSGA-II), which finds a set of potential solutions that are optimal for the combined objectives. Such solutions are known as Pareto optimal, where one of the objectives cannot be improved without negatively impacting at least one other objective. We further adjust NSGA-II to use the Factored Evolutionary Algorithm (FEA) framework, which decomposes the variables into separate, overlapping groups to increase exploration of the search space, as well as enabling the ability to parallelize computation.

Keywords.
Genetic Algorithm, Multi-Objective Optimization, Sustainability, Site-Specific Application, Economically Optimal Nitrogen Rate.

Introduction

Agricultural practices contributed 11% of greenhouse gases in 2020, and they are considered to be the single largest source of water pollution (USEPA, 2022). With the increase in concern regarding climate change and pollution, the question of how to increase sustainability in farming becomes more important. We propose using multi-objective optimization (MOO) to help address sustainability when creating variable rate fertilizer prescription maps.

In this paper we consider three different “objectives.” First, we seek to maximize net return for farmers. In addition, we also minimize two sustainability-focused objectives: reducing strain on applicators by reducing jumps between fertilizer rates and reducing the overall amount of fertilizer applied to minimize nitrogen pollution. When trying to optimize different problems (or objectives) simultaneously, this is called MOO. The challenge when attempting to optimize multiple objectives such as these arises from the fact tradeoffs exist. For example, in many cases minimizing the nitrogen across the field can result in a loss of net return due to a corresponding reduction in yield.

In previous research, we created experimental prescription maps for on-farm trials optimizing two objectives: maximizing stratification and minimizing jumps (Peerlinck et al., 2018). So instead of maximizing net return, we maximized stratification of fertilizer based on a field’s previous years’ yield content. We further extended these experiments by including the second sustainability-focused objective that reduces fertilizer. In these experiments, we found that we could reduce fertilizer by 5 to 10% without significantly influencing yield production. Because of these promising results, we now consider creating optimal prescription maps using the same techniques.

The preliminary research reported here focuses on how prescription maps that are optimized for the two sustainability objectives impact net return. We do this by applying three different variations of a multi-objective evolutionary algorithm called Non-dominated Sorting Genetic Algorithm II (NSGA-II) (Deb et al., 2002) to a winter wheat field and evaluating whether the resulting prescription maps’ net returns are significantly different. Consistent with our previous experiments on trial design, we found that incorporating the two sustainability objectives resulted in no significant impact on net return.

The rest of this paper is structured as follows. In the Background section we will briefly explain important terms and concepts in MOO, as well as the different algorithms we used. Next, we briefly review related work to provide the reader with a sense of other work that has been performed similar to ours. We will then move on to present the results of our initial experiments. Finally, in our conclusion we will discuss our main findings and future research directions.

Background

Meta-heuristics and Optimization

In Machine Learning (ML), a common approach to solving optimization problems is by applying what are called meta-heuristic approaches. Meta-heuristics are a class of approximate search methods designed to tackle difficult optimization problems when classical strategies are not sufficiently effective or efficient (Osman & Kelly, 1996). Within the set of meta-heuristics there are two different algorithm classes: local search and population-based search (Ehrgott & Gandibleux, 1994). Population-based algorithms, which are where our interests lie, are often further divided into swarm-based algorithms (e.g., particle swarm optimization (Kennedy & Eberhart, 1995)) and evolutionary algorithms (e.g., genetic algorithm (Holland, 1994)). In this paper we use a multi-objective version of the genetic algorithm (GA), called the Non-dominated Sorting GA II (NSGA-II) (Deb et al., 2002).

There are two large classes of problems meta-heuristics are often applied to: combinatorial and continuous optimization problems. In combinatorial optimization problems, one attempts to find an optimal combination of a finite set of objects, whereas in continuous optimization, one attempts to find the best real-valued variables (Osman & Kelly, 1996). There are a variety of reasons why these classes of problems are difficult, including the existing of multiple local optima, ruggedness in the objective landscape, and what is known as the "curse of dimensionality."

Both combinatorial optimization and continuous optimization exist and problems in precision agriculture: coverage path planning (Valenta, et al., 2016) is an example of combinatorial optimization problems, whereas wireless sensor network design contains both combinatorial and continuous problems (Ibrahim & Alfa, 2017). We handle the creation of an optimal prescription map as a continuous optimization problem. This means we are assuming the amount of fertilizer to be applied to the field can take any real-valued number within the bounds we set (i.e., between 0 and 150 pounds. of nitrogen per acre). In contrast, when we created experimental trial designs, the fertilizer to be applied came from a set of fixed fertilizer rates to be applied, making it a combinatorial problem (Peerlinck, et al., 2019).

Multi-Objective Optimization

Irrespective of a problem's continuous or combinatorial nature, optimization can involve attempting to address more than one objective simultaneously, creating a situation where the objectives compete, creating the necessity for a tradeoff. There are several ways to approach optimizing multiple objectives. One is known as a transformative approach, meaning the different objectives are transformed into a single objective (Deb, 2014). There are two commonly used transformative techniques: using a weighted sum, which assigns importance to each objective and linearly combines the associated objective functions and weights; and the ε-constraint technique, which chooses one objective as the main objective and transforms the other objectives into constraints limited by a set ε value.

However, both of these techniques assume knowledge on which objective is considered more important and produce a single solution based on this preference. Because of this, meta-heuristics using the idea of Pareto optimality are often used to produce a set of solutions that are considered to be non-dominated (Deb, 2014). Pareto non-dominated solutions are solutions where no other options exist that improves the result for one objective without deteriorating another objective; otherwise, this new solution would become the non-dominated solution. A set of non-dominated solutions with respect to all of the objectives is known as the Pareto optimal set (also referred to as a Pareto optimal front or frontier).

Implemented Algorithms

NSGA-II is an elitist Genetic Algorithm (GA) that finds Pareto non-dominated solutions and uses a crowding distance measure to maintain diversity in subsequent generations (Deb, et al., 2002). As with the classic GA, an offspring population is created using crossover and mutation. Afterwards, the parent population and the offspring population are combined into a single population. The resulting population is then sorted based on the Pareto non-domination principle, and individuals are assigned to different non-domination sets based on the extent to which they dominate other solutions in the population. If the set of non-dominated solutions is larger than the the pre-specified fixed size of the population, a second elimination is performed based on crowding distance.

In addition to NSGA-II, we consider two variants that employ principles from co-evolution. These two additional methods still use NSGA-II as the base algorithm, but instead of using a single population of individuals, there are now multiple subpopulations. Specifically, in the base NSGA-

II method, a population is made up of individuals corresponding to complete prescription maps. In the two co-evolutionary variants, each subpopulation's individuals represent a only a subset of the full map. This idea was proposed to allow for easy parallelization, which can help speed up computation when dealing with large scale problems. The first co-evolutionary variant is known as the Co-operative Co-Evolutionary Algorithm (CCEA) and was defined by Potter and De Jong (2000). CCEA uses disjoint subpopulations, meaning that each subpopulation considers disjoint sets of variables of the full target solution. In other words, with respect to optimizing prescriptions, each cell of the field can only belong to one subpopulation. In our experiments, we create subpopulations based on the strips of cells in the field, as shown by the dark rectangle covering column 7 in Figure 1. Each subpopulation evolves separately before being recombined to form the final global solution (i.e., the solution representing all cells on the map). Note, however, that while the subpopulations are organized into strips, the target is to provide finer grained, site-specific prescriptions at the cell level.

The CCEA approach was extended to include overlapping subpopulations, known as the Factored Evolutionary Algorithm (FEA) (Strasser et al., 2017). The main extension corresponds to permitting a single variable to be represented by multiple subpopulations, thus introducing "overlap" between the subpopulations being optimized. This adjustment was proposed so parallelization techniques can still be applied to the subpopulations, while also trying to improve exploration of the search space. For this algorithm, we still create subpopulations corresponding to strips, as with CCEA, but we include overlap by choosing the start and end cells of consecutive strips to create a new subpopulation; an example of this is shown in Figure 1. The arrows in columns 4 and 5 indicate an example of the direction the applicator takes, based on this direction, the lightly shaded square covering the bottom of columns 6 and 7 shows the two end cells of strip 6 and the start cells of strip 7. An MOO implementation of FEA is presented in (Peerlinck & Sheppard, 2022).

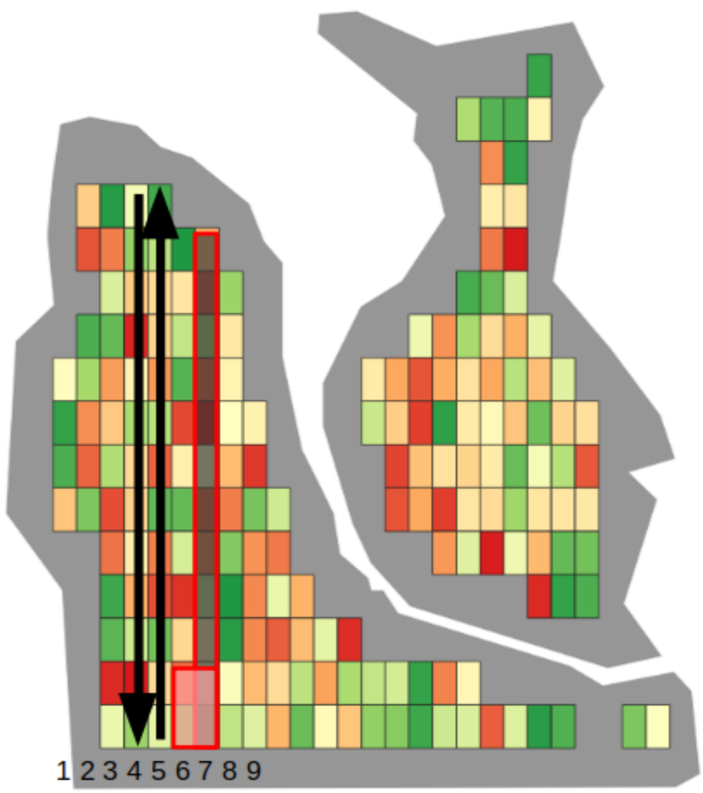

Fig 1. Example of an optimal prescription map. The arrows indicate the path the applicator takes along the plots. The shaded rectangle shows an example a subpopulation based on a strip of cells in the field. The shaded square is an example of a group of two ending and starting cells of consecutive strips to form overlapping subpopulations.

Preliminary Results

Experimental Design

As mentioned in the Introduction, we look at three different objectives for our experiments:

- Net return maximization
- Jump minimization
- Overall fertilizer minimization.

The Net Return (NR) is calculated as follows:

$$NR = Y * P - AA * CA - FC, \quad (1)$$

where *Y* is the expected crop yield, *P* is the crop selling price, *AA* is the 'as-applied' fertilizer rate, *CA* is the fertilizer cost, and *FC* reflects any fixed costs associated with production. The expected crop yield is predicted using a Random Forest (RF) model (Segal, 2004), whereas the crop prices and fertilizer cost are obtained from the US Department of Agriculture (USDA, 2022). Note that we are also in the process of developing more specialized yield prediction models.

Jump minimization addresses large changes in fertilizer rate application between consecutive cells in the field, since such large changes puts strain on the farming equipment. This can then lead to the farmer needing to repair or replace equipment more frequently, increasing cost and waste, thus resulting in negative ecological impacts. We address this objective by summing over the absolute differences in applied fertilizer between adjacent cells. Lastly, the fertilizer score is calculated by summing over the fertilizer to be applied to all of the cells.

To evaluate the impact of the different objectives, we chose three different non-dominated solutions produced by each algorithm, where the chosen solutions correspond to the extreme points for one of the objectives: jump score, net return, or fertilizer rate. Lastly, we picked a fourth solution to represent an equal balance between all three objectives. We do this by finding the centroid of the three extreme solutions, the non-dominated solution in the Pareto front closest to this centroid is then used as the fourth solution. We compared the four different types of prescription maps using an ANOVA test with $\alpha = 0.05$ to evaluate the impact of the different objectives on net return.

Results and Discussion

Based on the applied ANOVA, no significant difference in net return was found between the different prescription maps. This indicates that applying less fertilizer need not have a significant negative impact on farming profit. This is further confirmed when we inspect the net return values visually in Figure 2, where we can see that the difference in net return when focusing on different objectives is minimal for each algorithm. The union front combines all found non-dominated solutions for each of the algorithms. The best solutions for each of the aforementioned prescription map types are then select from this union front. The jump-focused prescription map found by the union front has a lower net return than the one found by F-NSGA-II because the jump score found NSGA-II was better, but that solution had a lower net return. Overall, the algorithms perform similarly in terms of net return, and they ran in a similar amount of time (2 days); however, we did not perform any parallelization on the CCEA and FEA implementations. By including parallelization of the subpopulations, we hope to reduce computational time, so we can provide results quickly to our end users. Furthermore, the lowest net return is found consistently when focusing on minimizing jumps. We believe this may be because the current net return calculation does not include the cost of equipment maintenance. If farmers could gather data on how large jump rates impact them economically, we could refine our net return calculation, and the difference in net return may be even less prominent.

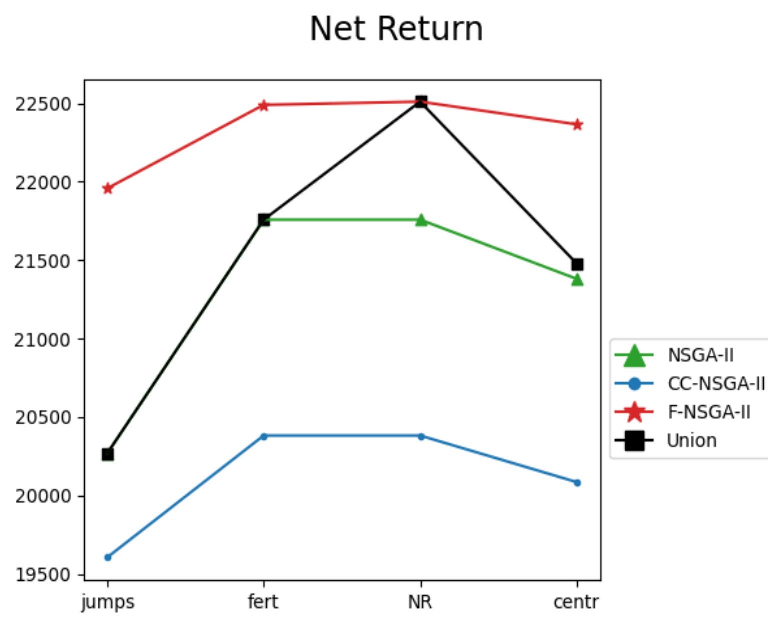

Fig 2. Net Return for four prescription maps focusing on different objectives.

Conclusion and Future Work

This work presents preliminary results optimizing fertilizer prescription maps that also consider sustainability focused objectives. We optimized the following three different objectives: maximizing net return, minimizing jumps between consecutive cells, and minimizing total fertilizer applied. We applied three different multi-objective optimization techniques to find a set of Pareto non-dominated solutions, where each solution represents a prescription map. We compared different prescription maps that are optimal for different objectives and found that there is no significant difference in net return between any of the different objectives. This is a promising indication that we can reduce fertilizer rate in variable rate application even more without impacting farming profit, thus improving environmental impact.

As next steps, we plan to apply our approach to more fields to confirm these preliminary findings. We would also like to add temporal objectives, such as minimizing variation in net return across several years, and including the impact climate change might have on crop response. Another goal is to investigate the effect of different yield prediction approaches when creating optimized prescription maps. In other words, how much influence does accurate yield prediction have on prescribing the correct fertilizer rate, or is it more important to use a model that accurately describes the shape of the yield response curve? Lastly, we would like to parallelize the subpopulation-based approaches to help reduce computational time.

Acknowledgements

This project was funded, in part, by MREDI grant 51040-MUSR12015-02, NSF grant 1658971, and USDA grant NR213A750013G021.

References

Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multi-objective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation, 6(2), 182–197.

Deb, K. (2014). Multi-objective optimization. In Search methodologies, 403–449. Springer.

Ehrgott, M. & Gandibleux, X. (2000). A survey and annotated bibliography of multiobjective combinatorial optimization. OR-Spektrum, 22(4), 425–460.

Holland, J. H. (1992). Genetic Algorithms. Scientific American, 267(1), 66–73.

Ibrahim, A., & Alfa, A. (2017). Optimization Techniques for Design Problems in Selected Areas in WSNs: A Tutorial. Sensors, 17(8). MDPI AG.

Kennedy, J., & Eberhart, R. (1995, November). Particle swarm optimization. In Proceedings of ICNN'95-international conference on neural networks (Vol. 4, pp. 1942-1948). IEEE.

Osman, I.H. & Kelly, J.P. (1996). Meta-Heuristics: An Overview. In Meta-Heuristics: Theory & Applications 1 – 21. Norwell, MA: Kluwer Academic Publishers.

Peerlinck, A., Sheppard, J., Pastorino, J., & Maxwell, B. (2019, June). Optimal Design of Experiments for precision agriculture using a genetic algorithm. In 2019 IEEE Congress on Evolutionary Computation (CEC) (pp. 1838-1845). IEEE.

Peerlinck, A. & Sheppard, J. (2022). Multi-Objective Factored Evolutionary Optimization and the Multi-Objective Knapsack Problem. In 2022 IEEE World Congress on Computational Intelligence.

Potter, M.A., Jong, K.A.D. (2000). Cooperative coevolution: An architecture for evolving coadapted subcomponents. In Evolutionary Computation 8(1), 1–29.

Segal, M. R. (2004). Machine learning benchmarks and random forest regression.

Strasser, S., Sheppard, J., Fortier, N., & Goodman, R. (2017). Factored evolutionary algorithms. In IEEE Transactions on Evolutionary Computation, 21(2), 281–293.

United States Department of Agriculture (USDA). (2022, January). Agricultural prices. https://usda.library.cornell.edu/concern/publications/c821gj76b?locale=en

United States Environmental Protection Agency (USEPA). (2022, April). Sources of Greenhouse Gas Emissions. https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions

Valente, J., Del Cerro, J., Barrientos, A., & Sanz, D. (2013). Aerial coverage optimization in precision agriculture management: A musical harmony inspired approach. In Computers and Electronics in Agriculture, 99, 153 – 159.