Addressing Sustainability in Precision Agriculture via Multi-Objective Factored Evolutionary Algorithms*

Amy Peerlinck & John Sheppard

Montana State University, Bozeman, MT, USA
{first}.{last}@montana.edu

Abstract. Precision agriculture is a research area that uses technology from engineering and computer science to improve all aspects of agriculture, including but not limited to crop health, irrigation, and fertilizer application. In agriculture, questions of sustainability often arise: How do we minimize environmental impact while simultaneously helping farmers maximize their net return? In this paper, we present a method to optimize crop yield production in winter wheat, with the goal of seeking to increase farmers' production. However, only focusing on optimizing production can lead to poor sustainability if an unnecessary amount of fertilizer is applied or farming equipment is put under undo stress. We therefore seek to address these impacts on sustainability by including objectives that directly address these concerns. Our method utilizes a new approach to solve multi-objective optimization that uses overlapping subpopulations, known as a Multi-Objective Factored Evolutionary Algorithm. Our results indicate that including overlapping subpopulations in the multi-objective optimization context is beneficial for exploration of the objective space. Our results also indicate that including these sustainability-driven objectives does not significantly impact net return or yield.

1 Introduction

Sustainable Agriculture has been defined by the U.S. Congress in the 1990 Farm Bill as "an integrated system of plant and animal production practices having a site-specific application that will, over the long term: Satisfy human food and fiber needs; Enhance environmental quality and the natural resource base upon which the agricultural economy depends; Make the most efficient use of non-renewable resources and on-farm resources and integrate, where appropriate, natural biological cycles and controls; Sustain the economic viability of farm operations; Enhance the quality of life for farmers and society as a whole [30]." Based on this definition, we investigate the effect of adding two sustainability-focused objectives to the problem of optimizing net return for winter wheat production, creating a multi-objective optimization (MOO) problem.

* Supported by NSF grant 1658971 and USDA Grant NR213A750013G021.

In our work, we focus on defining "prescription maps" that dictate levels of nitrogen fertilizer to be applied on a field planted with winter wheat. We address two problems: first is generating the experimental prescription maps, which are used to gather data on the field to determine the nitrogen response of the crop, and second is generating prescription maps based on estimating crop response to determine how much fertilizer to apply to maximize net return. Each problem has a unique objective: maximizing stratification and maximizing net return. However, they both have potential sustainability issues, which we address by considering two additional objectives: minimizing overall fertilizer applied to mitigate environmental impact, and minimizing jumps in variable rate application to reduce the economic impacts that can result from undue strain on equipment. Due to the resulting multi-objective nature of these problems, meta-heuristic approaches are often employed to explore possible solutions along the different objective axes [15]. This results in searching for a set of "non-dominated" solutions, i.e., solutions where no other options exist that improves the result for one objective without deteriorating another objective. [26].

In this paper, we consider three different MOO algorithms. First we apply the Non-Dominated Sorting Genetic Algorithm II (NSGA-II) [8] and the Co-operative Co-evolutionary NSGA (CC-NSGA-II) [19] to validate our approach. Then we extend the Factored Evolutionary Algorithm (FEA) [27] to create a multi-objective implementation which we call MO-FEA [21], where the specific implementation using NSGA-II is denoted F-NSGA-II. Then we apply all three approaches to the two problems in Precision Agriculture (PA): experimental and optimal prescription map generation.

2 Background

2.1 Prescription Maps

As discussed previously, our interesting in this paper is in developing prescription maps (experimental and optimal—Fig. 1) that optimize objectives balancing net return and sustainability. Experimental prescription maps aim to spread pre-defined fertilizer rates evenly across a field to determine crop response based on different input (e.g., nitrogen) levels [22]. Possible fertilizer rates are pre-defined to ensure that the different levels are represented evenly across the field. In our experiments, we focus on stratifying the rates based on previous years' yield information. In other words, we look at which parts of the field had high, medium, and low yield the year before and distribute the fertilizer across these 3 levels, where the specific objective function is based on the work in [22].

Optimal prescription maps specify fertilizer rates to apply based on crop response and economic models to maximize expected net return. These maps depend upon the ability to predict based on the prescribed inputs, general field information, and satellite data such as the normalized difference vegetation index (NDVI). The classic way to approach yield prediction is to use linear regression; however, this approach is limited in its ability to represent the yield response curves. As a result, machine learning approaches such as Random Forests [16]

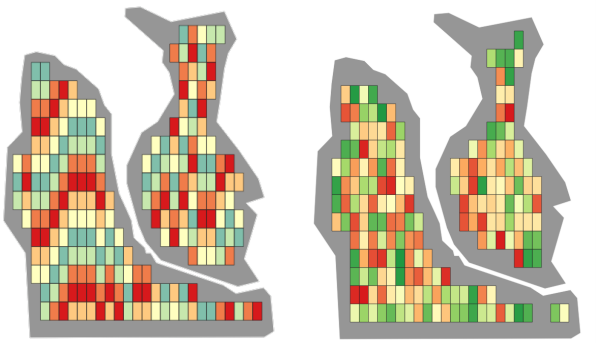

Fig. 1: Example of an actual experimental prescription map (left) and an actual optimal prescription map (right) for field Sec35Mid. Colors indicate fertilizer rates where red-to-green indicates increasing rates.

and Deep Learning [31] have become more popular. We use these models to predict the yield and use the result to determine the expected net return (NR):

$$NR = Y \times P - AA \times CA - FC, \tag{1}$$

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.

2.2 Multi-Objective Optimization

Multi-Objective Optimization is the process of optimizing multiple objective functions simultaneously [15]. Formally and without loss of generality, assume we wish to minimize k objectives. Then MOO consists of solving

$$\min_{\mathbf{x}\in\mathcal{X}} \mathbf{f}(\mathbf{x}) = \{f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_k(\mathbf{x})\},$$

with $k \geq 2$ objective functions $f_i : \mathbb{R}^n \to \mathbb{R}$ that have conflicting goals and $\mathbf{x} = [x_1, x_2, \ldots, x_n]^\top$ denotes the decision variables. Then for $f_i \in F^k$, F^k represents the objective space, and $\mathcal{X} \in \mathbb{R}^n$ represents the solution space.

One of the most approachable ways to perform MOO is to transform the set of objectives into a single objective [6]. The weighted sum approach does this by assigning weights to each objective function, where the weights sum to 1. The ϵ-constraint method transforms all objectives except one into constraints which are bound by a value $R\epsilon$. Then single-objective methods can be applied directly.

Such transformation approaches are limited by the fact that a user needs to decide which objective is more important. Alternatively, if a set of solutions is desired, the problem needs to be solved multiple times with different weights. To avoid such pitfalls, meta-heuristic approaches are often employed.

There are two different general classes of meta-heuristic approaches: local search and population-based search [9]. These meta-heuristics use the concept of Pareto optimality to return a set of potential solutions spread across the objective space.

When exploring algorithms that are generalized for any MOO problem, Pareto-based solutions are often favored [33]. More specifically, metaheuristic methods track the Pareto front of possible solutions along the different dimensions in the solution space [26]. Pareto dominance is used to determine solution quality. Formally, a point $\mathbf{x}^* \in \mathcal{X}$ is Pareto-dominant if $\forall f_i \in \mathbf{f}, \forall \mathbf{x} \in \mathcal{X}, \mathbf{x} \neq \mathbf{x}^*, f_i(\mathbf{x}^*) \leq f_i(\mathbf{x})$, and $\exists f_j \in \mathbf{f}, f_j(\mathbf{x}^*) < f_j(\mathbf{x})$. where $\mathbf{x} \prec \mathbf{y}$ denotes that $\mathbf{x}$ dominates $\mathbf{y}$. Then the Pareto optimal front (PF^*) is the set of points mapped from the set of Pareto optimal solutions onto the objective space F^k to form the boundary of the set of non-dominated solutions.

The Non-Dominated Sorting Genetic Algorithm (NSGA) was introduced by Srinivas *et al.* in 1994 [26], and improved in 2002 by Deb *et al.* to create NSGA-II [8]. It was then adjusted further to address MOO problems with more than three objectives to create NSGA-III [7]. However, we chose NSGA-II since it has been shown to be the better choice on 3-objective problems [13]. NSGA-II is an elitist GA that finds Pareto non-dominated solutions and uses a crowding distance measure to maintain diversity in subsequent generations. The parent population P_t and the offspring population Q_t are combined into one population $R_t = P_t \cup Q_t$. R_t is sorted based on the non-domination principle, and individuals are assigned to different non-domination sets based on how good the solution is. If an set of non-dominated solutions is larger than the remaining slots for the next population, a second elimination is performed based on crowding distance.

Cooperative co-evolutionary algorithms (CCEAs) were introduced by Potter and De Jong [23]. CCEAs are based on symbiotic relationships found in nature, where different species live together and improve each others' standard of life. To mimic this, a problem is divided into smaller components, each represented by a different population. In the first version of CCEA, single-objective problems with n dimensions are decomposed into n 1-dimensional subproblems, each with their own subpopulation. These subpopulations are then evolved separately and periodically recombined to form a complete solution. An individual's fitness is based, not only on how well it solves its own part of the problem, but also on its ability to cooperate with other partial solutions. CCEA is applied frequently to MOO, and the resulting algorithms are called co-operative co-evolutionary multi-objective optimisation algorithms (CCMOEA) [19]. Several studies found that a co-operative approach can be beneficial for finding a well spread out Pareto front when compared to the single population MOEA [12, 19].

3 Related Work

The question of sustainability in agriculture has existed since the start of precision agriculture as an area of study [2]. But addressing these sustainability issues in practice has proven more difficult. Variable Rate Application (VRA) involves technology that allows farmers to apply different input rates to differ-

ent parts of the field to control their production and reduce cost more precisely [25]. By using VRA, farmers are able to apply less fertilizer overall than if they were to apply a uniform rate across an entire field, thus also improving their sustainability practice [2].

Several studies have shown that VRA can help with sustainability [5, 32]; however, this is not always the case. Some studies found that using VRA increases the cost for the farmer [28]. To the best of our knowledge, three studies have applied MOO algorithms to VRA prescription maps to provide a set of potential prescriptions. In [22], a simple weighted-sum approach is applied to maximize stratification for experimental prescriptions while trying to minimize jumps in input levels between consecutive cells. Zheng *et al.* applied a multi-objective fireworks optimization algorithm for variable-rate fertilization [34] with the goal of finding the optimal fertilization for oil crops based on three objectives: yield, energy consumption, and spatial effects. An additional study used an economic optimization model to determine the fitness of fertilizer prescription maps and irrigation strategies for optimal crop yield in western Switzerland [18]. The authors integrated climate change into a genetic algorithm model and found that climate change increases income risk for farmers.

4 Optimizing Prescriptions with MO-FEA and NSGA-II

Classic CCEA only creates subpopulations that have distinct variables, i.e., there is no overlap between subpopulations. Strasser *et al.* proposed including overlap in subpopulations to create the Factored Evolutionary Algorithm (FEA). FEA defines a factor architecture (FA) to decompose the set of variables into groups, in a way that permits overlap. As such, FEA includes CCEA as a special case [27]. Due to the possibility of overlap, FEA combines principles from both cooperative and competitive co-evolution, and FEA with overlapping FA's has been shown to perform well on combinatorial optimization problems such as NK-landscapes [27] and Bayesian network abductive inference [11]. We extended FEA to handle MOO based problems in a way that still supports existing single-population MOO algorithms to work within the framework [21].

In MO-FEA, in addition to maintaining an archive of non-dominated solutions $\mathcal{N}$, we keep a set of global solutions $\mathbf{G}$ that are non-dominated for that generation, i.e., $\mathbf{G}$ is replaced each generation based on the non-dominated solutions found by the current subpopulations. Initially, each subpopulation is assigned the same global solution $g \in \mathbf{G}$ for evaluating the fitness of its individuals. As the algorithm progresses, a random global solution $g_r \in \mathbf{G}$ is chosen for each subpopulation without replacement to ensure that as many of the found non-dominated solutions are represented across the subpopulations.

In this work, we assume a field is divided into m cells, each of which will receive a prescription of the amount of nitrogen to apply. To create an experimental prescription we use a set of fertilizer levels to apply to the field (e.g., $F = \{0, 20, 40, 60, 80, 120\}$). This means we are solving a combinatorial optimization problem. On the other hand, the optimized prescription specify fertilizer rate

Algorithm 1: Factored NSGA-II

input : number of cells m, fertilizer values F, FEA generations it_{FEA}, population size n, NSGA-II iterations it_{NSGA}
initialize: individual $X \leftarrow \{x_0, \ldots, x_m\}; x_i \in F$, global solution set $\mathbf{X} \leftarrow \{X\}$, subpopulations $\leftarrow \{s_j \subset X\}$, non-dominated archive $\mathcal{N} \leftarrow \{\}$

1 **while** *FEA generations* $< it_{FEA}$ **do**
2 **for** *each* $s_j \in$ *subpopulations* **do**
3 $N''_{s_j} \leftarrow$ NSGA-II(s, n, it_{NSGA}, X)
4 **end**
5 $\mathcal{N}' \leftarrow \{\}$
6 **for** *each variable* x_i **do**
7 $X \leftarrow \text{random}(\mathbf{X})$
8 **for** *each* s_j *where* $x_i \in s_j$ **do**
9 $n'' \leftarrow \arg\min_{\text{crowding-distance}}(N''_{s_j})$
10 $\mathcal{N}' \leftarrow \mathcal{N}' \cup n''$
11 $X(i) \leftarrow n''_i$
12 $\mathcal{N}' \leftarrow \mathcal{N}' \cup X$
13 **end**
14 **end**
15 $\mathcal{N}' \leftarrow \text{non-dominated}(\mathcal{N}')$
16 $\mathbf{X} \leftarrow \mathcal{N}'$
17 $\mathcal{N} \leftarrow \mathcal{N} \cup \mathcal{N}'$
18 **for** *each* $s_j \in$ *subpopulations* **do**
19 $X \leftarrow \text{random}(\mathbf{X})$
20 $s_j(X) \leftarrow X$
21 $\arg\max_{\text{crowding-distance}}(\{p_0, \ldots, p_m\} \in s_j) \leftarrow X$
22 **end**
23 $\mathcal{N} \leftarrow \text{non-dominated}(\mathcal{N})$
24 **end**
25 **return** $\mathcal{N}$

on a continuum with a lower and upper bound. Thus we are solving a continuous optimization problem. In our experiments, we set the lower and upper limit of the fertilizer rate to $F = \{0, \ldots, 150\}$, since the farmers we work with impose an upper limit of 150 lbs nitrogen/acre.

Algorithm 1 shows the basic operation of F-NSGA-II, which is our implementation of MO-FEA where each subpopulation is optimized using NSGA-II [8]. During the "Compete" step (lines 5–17), overlapping subpopulations use the non-dominated solution in $\mathcal{N}''$ with the best crowding distance to represent the current decision variable. Each potential solution for every decision variable is then saved in a temporary solution set $\mathcal{N}'$. In addition, a randomly chosen non-dominated solution from the subpopulation is added to $\mathcal{N}'$ to increase exploration. $\mathcal{N}'$ is evaluated for non-dominance, only keeping Pareto optimal solutions, and the resulting $\mathcal{N}'$ forms the new set of global solutions $\mathbf{G}$. Then $\mathcal{N}'$ is added to $\mathcal{N}$, which is re-evaluated for non-dominance. Then $\mathbf{G}$ is shared across

the subpopulations, where the worst solution in the subpopulation is replaced by the chosen g. The fitness scores are then updated, completing the "Share" step (lines 18–22). This constitutes one iteration of MO-FEA.

5 General Experimental Approach

We pose the following hypothesis: Including sustainability-focused objectives to minimize overall fertilizer rate and reduce impact on farm equipment does not significantly degrade Montana winter wheat harvest profit. To evaluate our hypothesis we examine the cropping of winter wheat using both experimental and optimal prescription maps as applied to Montana fields using three MOO-algorithms, denoted NSGA-II (single population), CC-NSGA-II (cooperative co-evolutionary), and F-NSGA-II (factored).

5.1 Factor Architecture

The factor architecture for MO-FEA and CCMOEA is based on the length of a single strip, i.e., one group includes the cells from one end of the field to the opposite end where the applicator has to turn around. This gives us the distinct groupings as used in CCMOEA. To introduce overlap, we calculate a number p by multiplying the number of cells in each strip by 0.2 and rounding the resulting number upward such that $p \leq 1$. A new group is then created by combining the last p cells of a strip and the first p cells of the subsequent strip.

5.2 Objective Functions

To generate nitrogen prescriptions, we optimize the following objectives:

1. Maximize net return or stratification (see Section 2.1)
2. Minimize application level changes between adjacent cells
3. Minimize overall fertilizer rate

For the second objective, large jumps in fertilizer rate between consecutive cells puts strain on the farming equipment, increasing wear and tear. In turn, this leads to the farmer having to repair or replace equipment more frequently, increasing cost and waste, which has negative ecological impacts. To address this, we adjust the jump score presented in [22] to handle continuous fertilizer rates as follows: $Fitn_{jumps} = \sum_{i=1}^{c-1} |F(map_i) - F(map_{i+1})|$, where $F(map_i)$ is the fertilizer rate for the ith cell on the field. The jump score now sums over the absolute difference in applied fertilizer between adjacent cells determined by an "as-applied" map.

Lastly, we want to mitigate the effect fertilizer, such as synthetic nitrogen, has on the environment by reducing the overall amount of fertilizer applied to a field. This reduces pollution of the atmosphere by limiting greenhouse gas emission and can help avoid polluting waterways [17], which can result in a loss of drinkable water and the death of aquatic life. The overall fertilizer application is calculated by summing the fertilizer prescribed in each cell: $Fitn_{fert} = \sum_{i=1}^{c} F(map_i)$.

5.3 Evaluation Metrics

The hypervolume indicator (HV) is one of the most commonly used evaluation metrics in MOO [3]. Its popularity is partially because the only information needed to calculate the HV of a Pareto Front approximation is a reference point. This is in contrast to measures such as Generational Distance, which requires the true Pareto Front to be known. Since we do not know the true Pareto Front for our problems, the HV is a natural choice to gain insight in the size of the covered objective space [35]. In these experiments, we min-max normalize the objective scores and use a negative net return to support minimization. Given this, we know the worst possible solution in the objective space is $\{1, 1, 1\}$, which we use for computing HV. To assess the diversity of the Pareto Front approximations, we use the spread indicator SI [14], corresponding to the sum of the width for each objective, indicating how wide the solutions are spread across the objective space. For each algorithm, the average HV and SI were calculated across 5 runs. An ANOVA test with $\alpha = 0.05$ was performed to assess statistical significance across the algorithms, followed by a pairwise t-test with $p = 0.005$.

Finally, to compare two Pareto fronts, denoted as $\mathbf{X}'$ and $\mathbf{X}''$, we use the coverage C of the fronts as presented in [36]. We define the total non-dominated set, or union front, $\mathbf{X}^*$ to be the result of combining the fronts from the three algorithms' representative runs g as

$$\mathbf{X}^* = \text{nondom}\left(\bigcup_i^g \mathbf{X'}_i\right).$$

$\mathbf{X}^*$ can then be used to calculate what percentage of the original non-dominated set is included in $\mathbf{X}^*$: $AC(\mathbf{X}', \mathbf{X}^*)$. We then find the relative coverage of the non-dominated sets when compared to the union front, calculated as

$$AC(\mathbf{X}', \mathbf{X}^*) = \frac{|\{\mathbf{x}' \in \mathbf{X}' : \exists \mathbf{x}^* \in \mathbf{X}^* : \mathbf{x}^* \preceq \mathbf{x}'\}|}{|\mathbf{X}^*|}.$$

The union front's HV and SI are also calculated and compared.

5.4 Prescription Evaluation

For both experiments, we choose four different non-dominated solutions from the approximate Pareto Front created by each algorithm for each of the five runs of the algorithms. These solutions are based on the three extreme points in the Pareto Optimal set: minimum jump score, maximum net return, and minimum fertilizer rate. The centroid for these three solutions, $\mathbf{x}_c$, is found as follows,

$$\mathbf{x}_c^j = \frac{1}{3}\sum_{i=1}^{3} \mathbf{x}_i^j, \ \forall j \in k,$$

where k represents the objectives. The non-dominated solution closest to this centroid (based on the Lebesgue measure [4]), is used as the fourth solution. We compare the four different types of prescription maps using an ANOVA test ($\alpha = 0.05$) to evaluate the impact of different objectives on net return.

6 Experimental Prescription Results

6.1 Parameters and Data

For the experimental prescriptions, each of the algorithms was set to terminate after the non-dominated archive did not change for 5 iterations. We chose this approach since farmers are able to generate their own experimental trials, and they often create these trials shortly before they need to fertilize. Thus, we wanted to mimic the shortened runtime desired by the farmers. Mutation rate and crossover rate were set to 0.1 and 0.9 respectively, based on results from [22]. Note that we used swap mutation, and the parents for crossover are selected using tournament selection with tournament size 5. The remaining parameters are the population sizes for all three algorithms and the number of iterations NSGA-II needs to be run on the subpopulations for F-NSGA-II and CC-NSGA-II. To determine these parameters, a grid search was performed. Four different population sizes were considered, $\{100, 200, 500, 800\}$, and three different iteration limits, $\{50, 100, 200\}$. Based on the results of the grid search, a population size of 500 was chosen for all algorithms, and an iteration limit of 100 was chosen for both CC-NSGA-II and F-NSGA-II.

For our experimental maps, we collected data on three fields from two farms. We use the farmer designations for these fields (i.e., Henrys, Sec35Mid, and Sec35West). Previously, we trained a convolutional neural network (CNN) based on prior experiments to predict yield from the wheat harvested on these fields [1]. We use this CNN to predict yield based on the fertilizer applied. We create an initial, random prescription based on the field boundary. The farmer provides information on the width of their fertilizer application equipment and which fertilizer rates to apply across the field. For our experiments, the cell sizes for the fields are 120 ft by 300 ft. For all prescriptions, six different fertilizer rates were specified in pounds per acre: $F = \{20, 40, 60, 80, 100, 120\}$. Once the initial grid was created, the cells were ordered based on the "as-applied" route the farmer takes across the field to apply fertilizer.

6.2 Results and Discussion

We found significant differences ($\alpha = 0.05$) between yield predicted for each of the fields, as well as the different algorithms for each field. However, no significant difference was found between the results for different objectives, confirming our hypothesis that ethical objectives do not impact yield. For each field, the yield predictions for a specific prescription are averaged to create Fig. 2. A summary for each field of the hypervolume, spread indicator, and adjust coverage for each algorithm's non-dominated sets averaged over 5 runs, as well as the unionized non-dominated set, are given in Table 1. Adjusted coverage uses a randomly selected run for each algorithm to avoid bias. This is repeated five times and the five separate calculations are averaged to get the final adjusted coverage score.

The prescription maps across all three algorithms, as well as the union front, produce consistent yield predictions with small fluctuations between the different

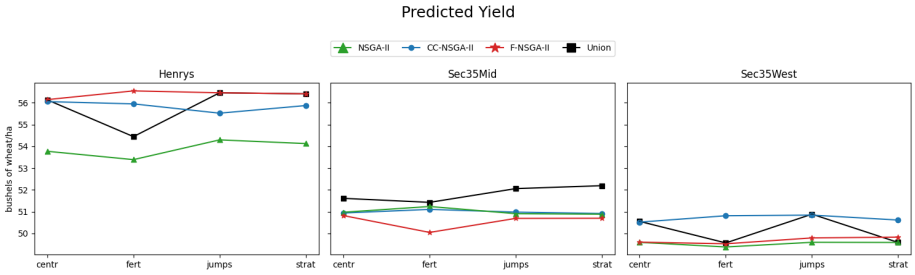

Fig. 2: Yield prediction results averaged across the entire field.

Table 1: Experimental prescriptions: Hypervolume (HV), spread (SI), and adjusted coverage (AC) results for each algorithm and the union front.

		NSGA-II	CC-NSGA-II	F-NSGA-II	Union Front
Henry's	**HV**	0.463	0.387	0.498	0.589
	SI	0.467	0.329	0.856	1.112
	AC	38.9%	11.3%	49.8%	100%
Sec35Mid	**HV**	0.465	0.397	0.504	0.593
	SI	0.440	0.289	0.824	1.011
	AC	30.2%	9.5%	60.3%	100%
Sec35West	**HV**	0.469	0.396	0.474	0.578
	SI	0.615	0.468	1.184	1.619
	AC	27.7%	26.2%	46.1%	100%

objectives, as can be seen in Fig. 2. The statistical results confirm what can be assessed visually: there is no significant difference between the predicted yield values across the different objectives, including those for the union front.

When looking at coverage between algorithms (Table 1), we can see that F-NSGA-II seems to cover more non-dominated solutions in the objective space than the other algorithms since it contributes the largest percentage of solutions to the union front. NSGA-II also makes large contributions to the union front, while CC-NSGA-II has the smallest contribution for all results. This could potentially be explained by the use of disjoint subpopulations in CCEA, since its disjoint nature means that a part of the solution space may be left unexplored. On the other hand, FEA uses the overlap to find more diverse solutions across the subpopulations, not only by saving a non-dominated solution from each subpopulation, but through the replacement of single variables in the global solution as well. The hypervolume and spread indicator results further confirm our hypothesis that F-NSGA-II explores more of the objective space. Across all fields, F-NSGA-II has the largest spread and HV for its approximate Pareto front. According to a pairwise t-test, the HV results were found to be significantly different for all three fields; however, no significant difference was found for the spread indicator results for Henrys and Sec35Mid between F-NSGA-II

and NSGA-II. When comparing the algorithms' results to the union front, we see that the union front HV and spread are not much larger than those found by F-NSGA-II, which is in line with the aforementioned coverage results.

7 Optimal Prescription Results

7.1 Parameters and Data

As in our previous experiments, mutation rate and crossover rate were set to 0.1 and 0.9 respectively. Swap mutation was used, and the parents for crossover were selected using tournament selection with tournament size 5. The remaining parameters are the population sizes for all three algorithms and the number of iterations NSGA-II needs to be run on the subpopulations for F-NSGA-II and CC-NSGA-II. For these experiments, our stopping criterion is the number of fitness evaluations. This approach leads to an increase in runtime; however, since farmers are not yet able to create their own optimal prescription maps, we can generate these maps beforehand, negating the need for a reduced runtime. To achieve this, we set the number of generations and population size such that each algorithm has approximately the same number of function evaluations. For NSGA-II, this resulted in setting a population size of 500 and running the algorithm for 500 iterations, yielding $500 \times 500 = 250,000$ FEs. Our instance of CC-NSGA-II has 24 groups, rounded up to 25 for ease of calculation, resulting in CC-NSGA-II being run 10 times where each subpopulation is of size 50, and NSGA-II is run for 20 generations on each subpopulation: $10 \times 20 \times 50 \times 25 = 250,000$. We used the same logic for F-NSGA-II, where population size is decreased to 25 to accommodate the increase in the number of subpopulations.

For this experiment, we look at field Sec35Mid since it is a more complex field and time constraints prevented us from finishing our experiments on the other two fields. The cell sizes are adjusted to be smaller (120 by 240 feet) to provide more refined optimal fertilizer rates. Based on the provided field data, we trained a regression model for yield prediction, where yield is reported in number of bushels per acre. We used the trained model to predict yield for the prescribed fertilizer; the predicted yield can then be substituted into Equation 1. We used economic data provided by the US Department of Agriculture to determine crop price and fertilizer cost [29].

We used a Random Forest (RF) as a regression model to predict yield. The RF model [16] was implemented via Scikit Learn [20] and evaluated via 10-fold cross-validation on the available data set. The RF yielded a root mean squared error of 2.97. Note that we were not able to use the neural network applied with the experimental maps since those maps could be evaluated offline while these needed direct interfacing between the optimizer and the network. Currently, that interface is still under construction.

7.2 Results and Discussion

Applying an ANOVA test to the net return results for the different types of prescription maps confirms that there is no significant difference in net return.

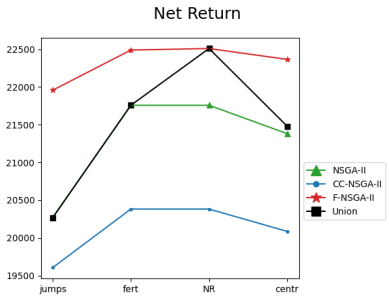

Fig. 3: Net return for the four different prescription maps.

Table 2: Experimental prescriptions: Hypervolume (HV), spread (SI), and adjusted coverage (AC) results for each algorithm and the union front on Sec35Mid.

	NSGA-II	**CC-NSGA-II**	**F-NSGA-II**	**Union Front**
HV	0.52	0.32	0.42	0.53
SI	1.67	1.65	1.68	1.68
AC	75%	0%	25%	100%

When we visually inspect the net return values in Figure 3, we can see that the difference in net return when focusing on different objectives is minimal for each algorithm. Interestingly, the lowest net return is found consistently when focusing on minimizing jumps. Currently, the net return calculation does not include the cost of wear on equipment. If farmers could gather data on how large jump rates impact them economically, we could refine our net return calculation.

When evaluating the algorithms' performance, we find there is no significant difference between the SI, but there is a significant difference for the HV ($p = 0.05$). Overall, NSGA-II performed better than F-NSGA-II and CC-NSGA-II for HV and coverage. F-NSGA-II, however, did consistently find solutions having a higher net return, indicating it explored a different part of the search space than the other two algorithms, which is reflected in the higher SI score.

Overall, single population NSGA-II performs better on the optimal prescription map problem; however, this could be due to the use of FE's as a stopping criterion. In [10], results indicate that using FE's may result in an unfair stopping condition. This idea is supported by our experimental map results, where we use convergence of the non-dominated archive as the stopping criterion. In this scenario, F-NSGA-II performs better; we believe this is because NSGA-II may be converging early on, before moving out of initial optima.

8 Conclusion

Multi-Objective Optimization provides a way for sustainability issues to be addressed when optimizing fertilizer prescriptions in precision agriculture. In this paper, we investigated adding two sustainability-focused objectives to existing precision agriculture problems, that of creating experimental and optimal fertilizer prescription maps. For both problems three competing objectives were optimized: the base objective, stratification or net return maximization respectively, and two sustainability objectives, fertilizer rate jump minimization, and overall fertilizer rate minimization. We applied three different MOO algorithms, NSGA-II, CC-NSGA-II, and F-NSGA-II, of which the latter is an adaptation of the Factored Evolutionary Algorithm in which overlapping subpopulations are used to find an approximate Pareto front. We found that all three MOO algorithms could find optimized prescription maps successfully, and that including these sustainability objectives had minimal impact on yield and net return. Based on these results, we confirmed our hypothesis that focusing on sustainability need not significantly influence net return, thus indicating a strong justification for modifying farming practices to incorporate such objectives, thereby reducing environmental impact. Furthermore, our results indicate that using overlapping subpopulations increases exploration of the objective space when compared to the single population and disjoint subpopulation alternatives.

9 Future Work

As next steps, we plan to investigate adding temporal objectives, such as minimizing variation in net return across several years, and including the impact climate change might have on crop response [18]. 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?

Results found by Strasser *et al.* indicate that the factor architecture could impact optimization [27]. However, work by Pryor *et al.* shows that the specific factor architecture may not matter as long as the overlap ensures that all variables are connected [24]. Based on these differing results, we plan to explore different factor architectures and their influence, not only regarding the problem of optimizing prescription maps, but by looking at different multi-objective benchmark problems such as DTLZ and the multi-objective knapsack problem.

Lastly, we would like to note that using function evaluations as a stopping criterion could impact results in an unfair way [10]. To this end, we intend to explore different ways to evaluate how long an algorithm should run, for example, based on the amount of change in the non-dominated archive as was done in the experimental design, or a lack of change in hypervolume. We would then evaluate the influence of different stopping criteria on runtime to enable a transition to farmers creating their own optimal prescription maps.

References

1. Anonymous, A.: Reduced-cost hyperspectral convolutional neural networks. Journal of Applied Remote Sensing **14**(3) (Sep 2020)
2. Bongiovanni, R., Lowenberg-Deboer, J.: Precision Agriculture and Sustainability. Precision Agriculture **5**(4), 359–387 (Aug 2004)
3. Bringmann, K., Friedrich, T.: Approximation quality of the hypervolume indicator. Artificial Intelligence **195**, 265–290 (Feb 2013)
4. Burk, F.: Lebesgue measure and integration: An introduction. John Wiley & Sons (2011)
5. De Koeijer, T., Wossink, G., Verhees, F.: Environmental and economic effects of spatial variability in cropping: nitrogen fertilization and site-specific management. In: The Economics of Agro-Chemicals, pp. 187–200 (2018)
6. Deb, K.: Multi-objective optimization. In: Search methodologies, pp. 403–449. Springer (2014)
7. Deb, K., Jain, H.: An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints. IEEE Trans. Evol. Comp. **18**(4), 577–601 (2013)
8. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multi-objective genetic algorithm: NSGA-II. IEEE Trans. Evol. Comp. **6**(2), 182–197 (2002)
9. Ehrgott, M., Gandibleux, X.: A survey and annotated bibliography of multiobjective combinatorial optimization. OR-Spektrum **22**(4), 425–460 (Nov 2000)
10. Engelbrecht, A.P.: Fitness function evaluations: A fair stopping condition? In: 2014 IEEE Symposium on Swarm Intelligence. pp. 1–8 (2014)
11. Fortier, N., Sheppard, J., Strasser, S.: Abductive inference in bayesian networks using distributed overlapping swarm intelligence. Soft Computing **19**(4), 981–1001 (April 2015)
12. Goh, C., Tan, K., Liu, D., Chiam, S.: A competitive and cooperative co-evolutionary approach to multi-objective particle swarm optimization algorithm design. European Journal of Operational Research **202**, 42—-54 (Apr 2010)
13. Ishibuchi, H., Imada, R., Setoguchi, Y., Nojima, Y.: Performance comparison of nsga-ii and nsga-iii on various many-objective test problems. In: 2016 IEEE Congress on Evolutionary Computation (CEC). pp. 3045–3052. IEEE (2016)
14. Ishibuchi, H., Shibata, Y.: Mating scheme for controlling the diversity-convergence balance for multiobjective optimization. In: GECCO. p. 1259–1271 (2004)
15. Ishibuchi, H., Tsukamoto, N., Nojima, Y.: Evolutionary many-objective optimization: A short review. In: IEEE CEC. pp. 2419–2426 (2008)
16. Jeong, J.H., Resop, J.P., Mueller, N.D., Fleisher, D.H., Yun, K., Butler, E.E., Timlin, D.J., Shim, K.M., Gerber, J.S., Reddy, V.R., et al.: Random forests for global and regional crop yield predictions. Public Library of Science **11**(6) (2016)
17. Kim, S., Dale, B.E.: Effects of nitrogen fertilizer application on greenhouse gas emissions and economics of corn production. Environmental science & technology **42**(16), 6028–6033 (2008)
18. Lehmann, N., Finger, R.: Optimizing whole-farm management considering price and climate risks. In: 123rd European Association of Agricultural Economists Seminar (February 2012)
19. Maneeratana, K., Boonlong, K., Chaiyaratana, N.: Multi-objective optimisation by co-operative co-evolution. In: Parallel Problem Solving from Nature. p. 772–781. Springer (2004)

20. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011)
21. Peerlinck, A., Sheppard, J., Maxwell, B.: Using deep learning in yield and protein prediction of winter wheat based on fertilization prescriptions in precision agriculture. In: International Conference on Precision Agriculture (2018)
22. Peerlinck, A., Sheppard, J., Pastorino, J., Maxwell, B.: Optimal design of experiments for precision agriculture using a genetic algorithm. In: IEEE CEC. pp. 1838–1845 (2019)
23. Potter, M.A., Jong, K.A.D.: Cooperative coevolution: An architecture for evolving coadapted subcomponents. Evolutionary Computation **8**(1), 1–29 (2000)
24. Pryor, E., Peerlinck, A., Sheppard, J.: A study in overlapping factor decomposition for cooperative co-evolution. In: 2021 IEEE Symposium Series on Computational Intelligence (SSCI). pp. 01–08. IEEE (2021)
25. Raun, W.R., Solie, J.B., Johnson, G.V., Stone, M.L., Mullen, R.W., Freeman, K.W., Thomason, W.E., Lukina, E.V.: Improving nitrogen use efficiency in cereal grain production with optical sensing and variable rate application. Agronomy Journal **94**(4), 815–820 (2002)
26. Srinivas, N., Deb, K.: Multiobjective optimization using nondominated sorting in genetic algorithms. Evolutionary computation **2**(3), 221–248 (1994)
27. Strasser, S., Sheppard, J., Fortier, N., Goodman, R.: Factored evolutionary algorithms. IEEE Trans. Evol. Comp. **21**(2), 281–293 (2017)
28. Thrikawala, S., Weersink, A., Fox, G., Kachanoski, G.: Economic feasibility of variable-rate technology for nitrogen on corn. American Journal of Agricultural Economics **81**(4), 914–927 (1999)
29. United States Department of Agriculture: Agricultural prices (Jan 2022), https://usda.library.cornell.edu/concern/publications/c821gj76b?locale=en
30. U.S. Congress: Agricultural research, extension, and teaching. In: U.S. Code Title 7, chap. 64. No. 3103, U.S. Government Publishing Office, Washington, DC, USA (2011), www.gpo.gov
31. Van Klompenburg, T., Kassahun, A., Catal, C.: Crop yield prediction using machine learning: A systematic literature review. Computers and Electronics in Agriculture **177** (2020)
32. Whitley, K.M., Davenport, J.R., Manley, S.R.: Differences in nitrate leaching under variable and conventional nitrogen fertilizer management in irrigated potato systems. In: Robert, P.C., Rust, R.H., Larson, W.E. (eds.) Proceedings of the 5th International Conference on Precision Agriculture. pp. 1–9. American Society of Agronomy, Madison, USA (2000)
33. Yu, X., Chen, W.N., Gu, T., Zhang, H., Yuan, H., Kwong, S., Zhang, J.: Set-based discrete particle swarm optimization based on decomposition for permutation-based multiobjective combinatorial optimization problems. IEEE Trans. Cybernetics **48**(7), 2139–2153 (2017)
34. Zheng, Y.J., Song, Q., Chen, S.Y.: Multiobjective fireworks optimization for variable-rate fertilization in oil crop production. Applied Soft Computing **13**(11), 4253–4263 (2013)
35. Zitzler, E.: Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications. Ph.D. thesis, Swiss Federal Institute of Technology (1999)
36. Zitzler, E., Deb, K., Thiele, L.: Comparison of multiobjective evolutionary algorithms: Empirical results. Evol. Comp. **8**(2), 173–195 (Jun 2000)