AN EVOLUTIONARY MULTI-OBJECTIVE APPROACH TO
SUSTAINABLE AGRICULTURAL WATER AND NUTRIENT OPTIMIZATION

By

Ian Meyer Kropp

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Biosystems Engineering – Master of Science

2018

ABSTRACT

AN EVOLUTIONARY MULTI-OBJECTIVE APPROACH TO
SUSTAINABLE AGRICULTURAL WATER AND NUTRIENT OPTIMIZATION

By

Ian Meyer Kropp

One of the main problems that society is facing in the 21st century is that agricultural production must keep pace with a rapidly increasing global population in an environmentally sustainable manner. One of the solutions to this global problem is a system approach through the application of optimization techniques to manage farm operations. However, unlike existing agricultural optimization research, this work seeks to optimize multiple agricultural objectives at once via multi-objective optimization techniques. Specifically, the algorithm Unified Non-dominated Sorting Genetic Algorithm-III (U-NSGA-III) searched for irrigation and nutrient management practices that minimized combinations of environmental objectives (e.g., total irrigation applied, total nitrogen leached) while maximizing crop yield for maize. During optimization, the crop model named the Decision Support System for Agrotechnology Transfer (DSSAT) calculated the yield and nitrogen leaching for each given management practices. This study also developed a novel bi-level optimization framework to improve the performance of the optimization algorithm, employing U-NSGA-III on the upper level and Monte Carlo optimization on the lower level. The multi-objective optimization framework resulted in groups of equally optimal solutions that each offered a unique trade-off among the objectives. As a result, producers can choose the one that best addresses their needs among these groups of solutions, known as Pareto fronts. In addition, the bi-level optimization framework further improved the number, performance, and diversity of solutions within the Pareto fronts.

To my family, to whom I owe where I am, and to my fiancé, who keeps me moving forward

ACKNOWLEDGMENTS

This work would not have been possible without my brilliant and dedicated professors, coworkers, friends and family. I stand on the shoulders of peers. In particular, I'd like to thank my major professor Dr. Pouyan Nejadhashemi for bringing together, mentoring, and efficiently running our diverse and prolific laboratory. Also, this research would not be possible without the pioneering research by Dr. Kalyanmoy Deb, as well as his patient teaching and mentoring. Finally, I thank Dr. Timothy Harrigan for keeping our research headed firmly toward practical real-world applications that improve our planet. This research also would not be possible without the direct contributions of Gerrit Hoogenboom, Mohammad Abouali, and Proteek Roy.

Equally important to this research are the intellectual and emotional support from my lab mates: Sebastian Hernadez, Melissa Rojas-Downing, Matthew Herman, Mohammad Abouali, Umesh Adhikari, Yirigui, Babak Saravi, Nathan Anthony.

Finally, I would not be where I am today if not for my friends, and family. In particular, my fiancée Kaitlyn, my siblings Julia and Emma and my parents are constant sources of inspiration and drive. Also, the members of an extremely long list of deep friends deserve significant credit. These include, but are not limited to, Andy Low, Barret Hoster, Emma Sanders, Levi Beach, Marisa Sandahl, and Mike Fox.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ALGORITHMS

KEY TO ABBREVIATIONS

BMP	Best management practices
CSM	crop simulation model
DE	Differential evolution
DSSAT	Decision Support System for Agrotechnology Transfer
EA	Evolutionary Algorithms
EMO	Evolutionary Multi-Objective Optimization
GA	Genetic algorithm
GP	Genetic programming
MO	Multi-Objective Optimization
MOCPSO	multi-objective chaos particle swarm optimization
MOFLP	Multi-objective Fuzzy Linear Programming
NP-hard	Non-deterministic Polynomial-time hard
NSGA-II	Non-dominated Sorting Genetic Algorithm II
NSGA-III	Non-dominated Sorting Genetic Algorithm III
PSO	Particle swarm optimization
SWAT	Soil and Water Assessment Tool
UF	University of Florida
U-NSGA-III	Unified Non-dominated Sorting Genetic Algorithm-III
WSM	Weighted Sum Method

1. INTRODUCTION

One of the major challenges that the world is facing in the coming decades is how to meet growing food demand without compromising the integrity of our environment (Mueller et al., 2012). An estimate suggests that global food production needs to be increased by 60-110% between 2005 and 2050 (Pradhan et al., 2015). Even so, by closing the yield gaps, which is the difference between attainable yield and actual yield in a region, most countries are expected to meet food self-sufficiency or to improve their current food self-sufficiency levels (Pradhan et al., 2015). Water and nutrient availability are the major production limiting abiotic factors in the regions where the yield gaps are high (Hengsdijk and Langeveld, 2009) and thus effective water and nutrient management plays a crucial role in food security by closing the yield gaps. In addition, optimizing water and nutrient management not only improves crop yield, but also reduces production cost, conserves resources, and protects the environment. However, the presence of multiple conflicting criteria, expensive simulation routines, nonlinearities in objective functions, and constraints make the optimization of such a system very difficult. Here, we are proposing to evaluate the performance of evolutionary multi-objective optimization methods as an alternative approach for addressing these types of socio-economic problems. To test this hypothesis, this thesis seeks to address the following research objectives through the utilization of evolutionary multi-objective optimization methods: 1) identification of the best irrigation practices to achieve high crop yields at minimum water usage, 2) identification of the best irrigation and nutrient management practices to achieve high crop yields at lowest environmental cost, 3) evaluate the importance of the number, time, and amount of irrigation and fertilizer applications on crop yields.

2. LITERATURE REVIEW

The application of optimization in agricultural intensification has a long and rich history. Researchers have applied classic single objective optimization techniques (e.g., linear programming, dynamic programming, and genetic algorithms) to a wide range of applications since the 1960s (Flinn and Musgrave, 1967). But within the last 25 years, agricultural engineers have embraced a new and powerful class of optimization algorithms know as multi-objective optimization (MO).

This literature review seeks to aggregate and analyze the current state of the art applications of MO algorithms within the concept of agricultural intensification. In this literature review we first introduce nutrient and water management in *Nutrient and Water Management in Agricultural Intensification*, and then we summarize currently ubiquitous applications for MO algorithms in agricultural intensification in the *Optimization Objectives* section. In the following section *Optimization Techniques*, the literature review then describes and analyzes optimization techniques within the literature and covers case studies of algorithms and their applications in agricultural intensification.

2.1 Nutrient and Water Management in Agricultural Intensification

As water and nutrients are the limiting abiotic factor in agricultural intensification, they are the decision variables in focus in this review. For water specifically, as the sustainability of agricultural water use is affected by competition from non-agricultural water use and climate change, there is an increasing interest in minimizing agricultural water use through improving water productivity (Morison et al., 2008). Water productivity is defined as crop yield per volume of water applied (Kijne et al., 2003). Water applied to the cropped fields can be lost either as a productive (transpiration) or as an unproductive (soil evaporation, infiltration, and runoff) water. Increasing water productivity entails increasing the productive water use while minimizing the unproductive water losses. By optimizing the irrigation schedule and increasing the water productivity, more amount of crop can be harvested with the same amount of irrigation water. In dry areas where cultivated land is limited by the lack of sufficient water, optimizing water productivity at a farm scale would help bring more area under cultivation by increasing water availability. In addition, optimizing the operation of regional water systems (e.g., irrigation networks, reservoirs) can increase the area of land under cultivation by assuaging water shortages in arid regions and regions with non-agricultural competition.

Similarly, fertilization is essential for increasing crop productivity; however, over-application or incorrect timing of fertilization may lead to contamination of surface and groundwater (Adesemoye et al., 2008). Focusing on a single nutrient, such as nitrogen, and its over-application causes nutrient imbalance, economic loss, and environmental pollution (Goulding et al., 2008), while under application leads to poor crop yield. Nutrient management involves managing the amount, source, timing and method of nutrient application to minimize nutrient loss and maximize plant uptake (Gaskin and Wilson, 2009). Nutrient management optimization synchronizes fertilization application with plant nutrient utilization, which maximizes crop yield and quality, increases profit, conserves resources and enhances soil quality and productivity. This is important to ensure long-term food security through a proper balance between increased food production, soil health and environmental quality (Lamessa, 2016).

Furthermore, water and nutrient management together have an even greater impact on agricultural intensification. Under a limited water supply, plant nutrient plays an important role in enhancing water productivity (Waraich et al., 2011). Under normal water supply condition, transpiration rate is increased by fertilization; however, under dry condition, fertilization has been found to result in depressed plant growth and higher seedling mortality rate (Li et al., 2009; Rahimi et al., 2013). Effective water management improves nutrient availability and helps the transformation of nutrients in the soil (Li et al., 2009). Hence, in the regions where both nutrient and water availability are constraints to crop yield, combined water and nutrient management is essential in increasing crop production. Therefore, optimizing nutrient and water management help improve food security by closing the yield gap, especially in the developing world. For example, for maize in Sub-Saharan Africa, closing the yield gap to 50% of the attainable yield can be achieved through nutrient management, but to close the yield gap to 75% of the attainable yield, simultaneous nutrient and water management is required (Mueller et al., 2012).

To computationally optimize agricultural intensification, it is necessary to develop efficient computation models for these agricultural systems. The arrival of physiologically based crop growth models allowed researchers to simulate plant growth and yield under varying irrigation and fertilizer supply. Some of the widely used model such as GOSSYM (Baker et al., 1983), CROPGRO (Boote et al., 1998), CERES-Maize (Jones et al., 1986), CERES-Wheat (Ritchie, 1985), SOYGRO (Wilkerson et al., 1983), PNUTGRO (Boote et al., 1992), AquaCrop (Steduto et al., 2009) and CropWat (Smith, 1992) have been used and improved in the past few decades. Hydrological models such as SWAT (Arnold et al., 2012, Neitsch et al., 2011), TOPMODEL (Kirkby, 1975), and MIKE SHE (DHI, 2003) allow researchers to evaluate the environmental impact, agricultural productivity and economic productivity of agricultural practices on entire regions.

2.2 Optimization Objectives

MO algorithms are currently applied to two broad categories: micro agricultural management and macro agricultural management. Micro agricultural management attempts to optimize the performance in a single agricultural unit (e.g., a single field or a single farm enterprise), where macro agricultural management attempts to optimize the performance multiple individual agricultural units at the watershed, county, or regional scales.

2.2.1 Micro Agricultural Management Applications

The less ubiquitous micro agricultural management applications of MO include irrigation management (Akbari et al., 2018; García-Vila et al., 2009), and crop planning (Groot et al., 2012; Mello Jr et al., 2013; Sarker and Ray, 2009). Even at a smaller scale, agricultural intensification problems still pose multiple conflicting objectives at the micro agricultural management scale and are therefore are ideal for MO techniques. For example, trying to minimize the levels of nitrogen loss in a given field will conflict with the overall yield of the crop (Hengsdijk and Langeveld, 2009).

2.2.2 Macro Agricultural Management Applications

There are abundant examples of MO being applied to macro agricultural management objectives, such nutrient tax policy (Whittaker et al., 2017), irrigation network operation (Ashofteh et al., 2015; Fernández García et al., 2014), land use (Groot et al., 2007), and regional crop planning (Sarker and Ray, 2009, 2005; Tan et al., 2017; Wang et al., 2012). MO is a popular choice in macro agricultural management because there are large numbers of conflicting objectives on a regional scale. For example, producers near a river may over apply nutrients to maximize their output, where regional governments may seek to minimize consequent algal blooms in a local reservoir. Or perhaps the consistent hydrologic head required for hydroelectric power generation is interrupted by high irrigation demands in the tropical dry season (Quinn et al., 2018). In addition to being multi-objective, macro agricultural management applications are ideal for evolutionary algorithms (section 2.3.1). Irrigation scheduling, for example, is an NP-hard problem (Anwar and Haq, 2013). The set of NP-hard problems is defined as the set of problems that have not been solved by algorithms in polynomial time, and are possibly be unsolvable in polynomial time (Cormen et al., 2009). The difficulty of such irrigation problems renders most simple brute force algorithms unrealistic, and evolutionary algorithms are therefore popular choices in the literature.

2.3 Optimization Techniques

2.3.1 Classical and Evolutionary Single Objective Optimization Approaches

In general terms, optimization algorithms search for optimal solutions within a decision variable space. Decision variable space is the space containing all the possible choices that a decision-maker can implement. One example decision space in agriculture would contain all the possible combinations of irrigation dates and amounts (between 0 and 50 mm) for a 120-day growing season. With 51^{120} or $8.09 \cdot 10^{204}$ different solutions, this decision space is too large for a human to reasonably evaluate in its entirety. Under these types of conditions, optimization algorithms can be useful tools. Optimization problems are typically defined by one or more decision variables (e.g., when to irrigate) and by one or more objective functions that numerically define the performance of a given solution (e.g., the seasonal yield for a single irrigation scheme).

Classical optimization techniques, for this paper, optimize only a single objective function and are fully deterministic (Deb, 2009). This class of algorithms typically has excellent performance with a certain subset of optimization problems (e.g., a differentiable, linear, or unimodal objective function), and includes, among others, Quasi-Newton, gradient descent, linear programming, and golden search section search (Deb, 2009). But outside their narrow scopes of high performance, classical optimization techniques struggle to search for optimal solutions in highly non-linear, non-differentiable, multi-modal, and/or multi-objective problems (Goldberg, 1989).

Evolutionary optimization techniques offer a less specialized, and more flexible approach to optimization. Where classical methods solve problems deterministically, evolutionary algorithms traverse search spaces with stochastic heuristics inspired by the phenomenon of evolution. Genetic algorithms (GA), a widely used subset of evolutionary algorithms, mimic evolution by creating an initial random "population" of solutions that evolve towards more and more ideal solutions after each generation (Gen and Cheng, 2000). Each individual of a population has a set of "genes" that represent the specific decision variables for that given solution, and each solution in a population is evaluated and ranked using a fitness function (i.e., objective function). The fittest solutions are paired with each other, and their genes are recombined into offspring solutions. Mutation operators create further diversity in the population. Ideally, the population as a whole will converge on an optimal solution, though there is no way to guarantee a solution is the true optimal solution. Similar to DNA, the genes can be coded as binary strings that can br crossed over with the genes of a mate solution (Holland, 1975). Alternatively, in real coded GA, genes can also be coded as arrays of real numbers and genes are crossed over using a process known as simulated binary crossover (Deb and Agrawal, 1995).

2.3.2 Multi-Objective Optimization

What differentiates single objective algorithms and MO algorithms is the number of and nature of objectives. Single objective algorithms on the one hand search for a solution that satisfies a single objective, and MO algorithms on the other hand search for solutions that satisfy multiple conflicting objectives. Conflicting objectives are objectives that cannot be satisfied with a single solution. For example, a hypothetical producer wants to optimize his or her urea application practices to a) maximize crop yield and b) minimize nitrogen application totals. The ideal solution that maximized crop yield would require generous amounts of urea, while the ideal solution that minimized total urea applied would use no urea at all. Therefore, this example has at least two equally optimal solutions: a solution that maximizes yield and a solution that minimizes total urea application. In other words, an optimization problem with two conflicting objectives will have a two-dimensional set of equally optimal solutions, and an optimization problem with *n* conflicting objectives would have an *n*-dimensional set of equally optimal solutions (Goldberg, 1989). These sets of solutions are known as Pareto fronts, and each member of the Pareto front are known as a non-dominated solution (Tamaki et al., 1996). Non-dominated solutions are solutions in a population that are not dominated by any other solution, where dominance is defined as outperforming a solution in every single objective. In summary, algorithms that can effectively search through multi-objective solution space are highly applicable to outstanding agricultural engineering problems with conflicting multiple objectives.

2.3.2.1 Classical Multi-Objective Approaches

During the naissance of the MO field, algorithms resolved conflicting objectives by reducing multiple objectives to a single objective. Once reduced to a single objective, a single objective optimization technique will find the optimum. These "classical" MO algorithms reduce the search space into a smaller region of the Pareto front. Subsequently, MO algorithms often allow for search within different regions of the Pareto front. The family of classical MO algorithms includes weighted sum, Tchebyshev (Miettinen, 2012), Benson's (Benson, 1978), and ε-constraint methods (Haimes, 1971).

Several papers in the literature employ the weighted sum approach. In its most common form, the weighted sum approach multiplies an objective vector *O* of *n* objectives by a weight vector *w*, and then sums the items of the product vector together into a single objective.

$$O_{overall} = \sum_{n}^{i=1} O_i w_i$$

With the multi-objective vector reduced to a single value, a classical single objective technique (e.g., GA or Linear Programming) will then use $O_{overall}$ as an objective function. The weight vector represents the importance given to each objective by a human decision-maker, and relatively higher weights endow a given objective more impact on the fitness of a solution.

There are several applications of the weighted sum method (WSM) in agricultural engineering. Nixon et al. (2001) applied the WSM to solve off-farm irrigation channel delivery schedules.

Nixon's framework maximizes "the number of orders that are scheduled to be delivered at the requested time" and minimizes "…variations in the channel flow rate." Behind the WSM lies a single objective genetic algorithm. Tan et al. (2017) reduced a multi-objective problem to a single objective fuzzy-robust linear programming problem, using relative membership grades as the weights. Sarker and Ray (2009) also used the WSM to validate the results of an evolutionary multi-objective optimization (EMO) algorithm, the Non-dominated Sorting Genetic Algorithm II (NSGA-II), in a crop planning problem. Using WSM is a common validation tool in EMO research (Deb, 2009).

The epsilon constraint method (Haimes, 1971) also converts a multi-objective problem into a single objective problem but instead uses constraints to reduce the number of objectives. In an n objective problem, $(n - 1)$ of the objectives are constrained to a single value, and the remaining objective is solved using a single objective method. Consoli et al. (2008) translated a two objective problem down into a single objective non-linear programming problem. The researchers, trying to 1) minimize irrigation deficit and 2) maximize net economic benefits, constrained the second objective function while minimizing the first objective function. Sarker and Ray (2009) used the epsilon constraint method to validate their NSGA-II crop planning optimization.

Other classical approaches include Genetic Programming, as used by Ashofteh et al. (2015) to minimize regional vulnerability to irrigation deficits and to maximize reservoir reliability. Two optimization scenarios, one in recent past and one in the near future, produced unique solutions to two unique climate scenarios.

2.3.2.2 Evolutionary Multi-Objective Optimization

During the 1990s and 2000s, EMO techniques grew in popularity. Evolutionary algorithms (EA) are powerful tools for solving MO problems because EAs search for populations of optimal solutions each generation. Therefore, a population-based approach allows an algorithm to search for an entire Pareto set of solutions.

2.3.2.2.1 Non-Dominated Sorting Evolutionary Multi-Objective Optimizations

In MO problems, non-dominated sorting is an effective means of ranking solutions in a population based on their convergence. Suggested by David Goldberg (1989), non-dominated sorting groups solutions into increasingly better non-dominated fronts (Figure 1) in terms of convergence. A non-dominated front contains solutions that are all non-dominated (i.e., no solution in the set dominates another solution within the set). The process starts by finding all of the non-dominated solutions in a population. These solutions become the first and highest ranked non-dominated front. The front is then removed from the rest of the population, and the second highest ranked front is identified. The process is repeated until all solutions are categorized into ranked non-dominated fronts. Once so ranked, an algorithm can quickly identify the dominance relationship between members in a population. For example, all solutions in the third highest ranked non-dominated front would automatically be chosen for the next generation over a solution from the fourth highest ranked non-dominated front.

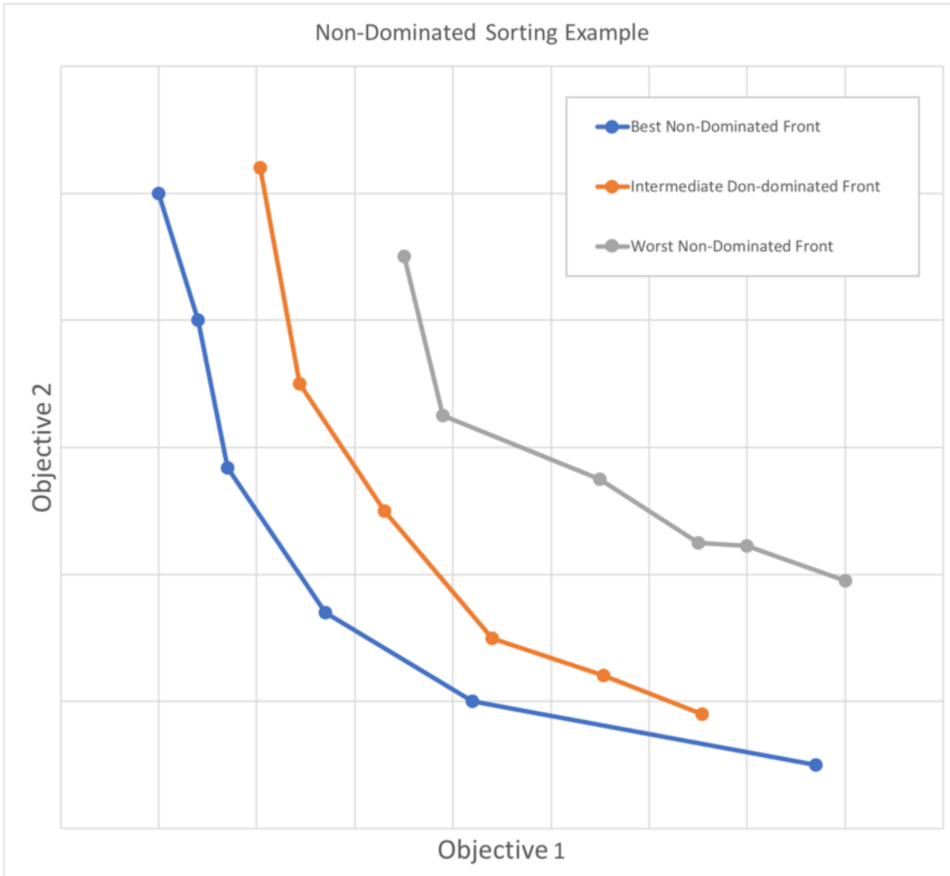

Figure 1. Example of non-dominated sorting in a two objective minimization problem

Agricultural researchers have employed members of the NSGA-II family of algorithms consistently since the mid-2000s. NSGA-II is a powerful, quick, and simple EMO algorithm developed by Deb et al. (Deb et al., 2002), which uses a combination of non-dominated sorting (Goldberg, 1989) and crowding distance to rank the fitness of a population of solutions. The algorithm balances the goals of convergence through non-dominated ranking, and then population diversity through crowding distance ranking. NSGA-II performs best with one and two objectives, decently with three objectives, and poorly four or more (or "many") objective problems. NSGA-II also implements the concept of elitism, in which the parents in one generation, instead of being removed from the population after their given generation, have the opportunity to be carried on to the following generation (Deb et al., 2002).

NSGA-II appears in a wide breadth of papers in the agricultural engineering literature. Its popularity mainly stems from its ease of use, simple parameters, and balance of population convergence and diversity. In two papers, Sarker and Ray (2009, 2005) developed and employed a variant on NSGA-II to optimize crop-planning practices. Both studies examined an agricultural region containing multiple farms and sought to assign crops to each farm optimally. In the first paper, the total economic investment in the region was minimized while the regional profit was maximized (Sarker and Ray, 2005). For the second paper, the total gross margin was maximized while the total cultivation cost was minimized (Sarker and Ray, 2009). Another research group, Darshana et al. (2012), attempted to maximize gross economic output and minimize water usage in a region in Ethiopia by changing the cultivars of three different farms in Ethiopia. NSGA-II optimized two objectives: an objective to maximize net benefits for farmers and an objective to minimize the water requirements from the crops. The research of Perea et al. (2016) optimizes pressurized irrigation, specifically sectoring operations. Using a customized version of NSGA-II, Perea et al. (2016) minimized the cost of running pumping stations while maximizing farmer's profit in a region in Spain. In another application of NSGA-II, Lalehzari et al. (2016) studied the optimal allocation of groundwater and surface water for deficit irrigation. Lalehzari et al. (2016) aimed to minimize the total water allocated and maximize the profit relative to the production costs. The authors then ran the same multi-objective optimization problem using particle swarm optimization. Most recently, Whittaker et al. (2017) developed a unique bi-level optimization routine to determine the optimal fertilizer tax on a watershed. Inspired by the Stackelberg game (Von Stackelberg, 2010), they first optimized the spatial distribution of a proposed fertilizer tax from the perspective of a policymaker, using NSGA-II to maximize the agricultural output of the region and to minimize the environmental impact. This is the upper level of optimization. In the lower level of optimization, farmers react to the fertilizer tax and optimize their profits with linear programming. The process then repeats with the upper level. A hybrid of chaos algorithm (Jiang and Weisun, 1998) and two MO algorithms (NSGA-II and MODE) appears in the work of Arunkumar and Jothiprakash (2017). They applied their modified algorithms to optimize crop planning in a multi-reservoir system spanning multiple basins and aim to maximize the net benefits and crop production.

There are many examples that use other algorithms based off of non-dominated sorting. ε-NSGA-II, a modified version of NSGA-II developed by Kollat and Reed (2006), incorporated the concepts of ε-dominance archiving into the NSGA-II. The goal of ε-dominance archiving is to perform a more uniform and spread out search within the objective space (Laumanns et al., 2002). ε-NSGA-II searches for watershed best management practices (BMP) in Liu et al. (2013). The objectives of the run are to minimize the cost of the BMP while maximizing the reduction in phosphorus load. The Soil and Water Assessment Tool (SWAT) model predicted the phosphorus load for a given BMP practice (Arnold et al., 2012, Neitsch et al., 2011). Zhang et al. (2017) also employ ε-NSGA-II to minimize agricultural water shortages while simultaneously optimizing other competition water requirements in a watershed and minimizing environmental impact. The algorithm developed by Groot et al. (2012) incorporates Non-dominated sorting into differential evolution to optimize overall farm management. The objectives included minimizing nitrogen leached/denitrified and labor requirements, as well as maximizing economic benefit and organic matter balance.

2.3.2.2.2 Other Evolutionary Multi-Objective Optimization Algorithms

There are other examples of EMOs applications in the agricultural intensification optimization literature. For example, differential evolution (DE) (Storn and Price, 1997) is a popular variant of evolutionary algorithms in the agricultural management optimization literature, and a multi-objective version of DE (Lampinen et al., 2000; Xue et al., 2003) appears in Groot et al. (2007).

The study by Groot et al. (2007) optimized land use and hedgerow placement with respects to area yield, biodiversity and nutrient loss.

Some papers use swarm-intelligence-based algorithms. First coined by Beni and Wang (1993), swarm intelligence mimics the behavior of swarms of autonomous biological agents (e.g., birds, ants, wolves). Wang et al. (2012) employed a variant swarm intelligence technique known as particle swarm optimization (PSO), which searches optimal solutions with a “swarm” of particles that act like a flock of birds (Eberhart and Kennedy, 1995). With each iteration, each “individual” of the swarm moves to a new position based on its current velocity, its personal best location at that time, and a global best location of the entire swarm. However, the variant of PSO, multi-objective chaos particle swarm optimization (MOCPSO), incorporates a dynamic weighted sum optimization within each particle of the swarm to optimize crop planning and water resources in a region in China. The objectives included maximizing the regional agricultural output, total grain yield, environmental benefit, and water efficiency.

Genetic Programming (GP) is another family of algorithms that follow evolutionary principles. Unlike GAs, which optimize binary strings representing possible solutions, GP optimizes computer programs (i.e., mathematical functions) in the form of parse trees (Koza, 1992). Ashofteh et al. (2015) employed a bi-objective MO version of GP to maximize the reliability of reservoir irrigation responses and minimize regional vulnerability to irrigation deficits.

Melody Search (Ashrafi and Dariane, 2013), a variant of Harmony Search (Geem et al., 2007) appears in the work of Karami and Dariane (2018). Both Melody and Harmony Search attempts to mimic how musicians improvise melodies within a musical ensemble. In Karami and Dariane, a Melody Search framework simultaneously optimizes multiple climate scenarios with respects to maximizing municipal, instream requirement, agricultural and hydropower reliability.

2.3.2.2.3 Many-Objective Optimization

Many objective problems are multi-objective algorithms that have more than three objectives. Many of the first generation of multi-objective algorithms perform poorly after three objectives, though a number of many objective algorithms have cropped up in the last ten years, such as MOEA/D (Zhang and Li, 2007), Borg (Hadka and Reed, 2013), and NSGA-III (Deb and Jain, 2014). Many-objective problems have not significantly appeared in the agricultural intensification optimization literature. Among the few are Gurav and Regulwar (2012), who used a Multi-objective Fuzzy Linear Programming algorithm (MOFLP) to solve a four objective irrigation planning problem. With MOFLP, irrigation planning in a region in India is optimized with respects to maximizing manure utilization, crop production, job creation, and the overall economic benefit. The research of Wang et al. (2012), mentioned earlier for their PSO algorithm MOCPSO, optimizes a four objective crop planning and water resource problem. In another publication, Karami and Dariane (2018) overcame the difficulties of many-objective optimization by combining Melody Search (Ashrafi and Dariane, 2013) with the concepts of social choice (Arrow, 1951; de Borda, 1781). With their hybrid optimization algorithm, Karami and Dariane (2018) simultaneously maximize reliability for regional municipalities, instream conditions, agricultural production, and hydropower operations under four different climate scenarios. Four NSGA-III runs, one for each climate scenario, is also used to optimize the aforementioned four objectives and to validate the results of the combined melody search and social choice algorithm. Zhang et al. (2017) employed ε-NSGA-II (Kollat and Reed, 2006) to solve a five objective regional water planning problem, which optimized against the water demands of businesses, agriculture, and the environment.

2.4 Literature gaps

There are several gaps in the agricultural intensification optimization literature. Firstly, and to the best of our knowledge, there are no micro agricultural management applications that combine both nutrient and water management. As mentioned earlier, simultaneously optimizing nutrient management and water management outperforms optimizing each objective individually (Waraich et al., 2011). Therefore, micro-managing nutrient and irrigation applications on farms would move forward the use of state-of-the-art techniques in agricultural intensification.

Furthermore, optimizing only irrigation and nutrient applications would ignore the basic requirement to make any management practice economically viable. Therefore, simultaneously optimizing irrigation management, nutrient management, and crop yield would balance three highly significant objectives in agricultural intensification. Finally, including an objective for environmental impact would add a powerful perspective on how certain agricultural practices affect the environment as a whole. Adding environmental objectives would determine whether or not a practice could be sustainably applied throughout regions.

Also, there is a lack of good applications of many objective algorithms in the agricultural intensification optimization literature. Agriculture, with its many conflicting objectives, would be an excellent case study for many objective problems. Agricultural systems are highly complex and non-linear and could serve to better define the strengths and weakness of the current state of the art many objective algorithms.

3. MATERIALS AND METHODS

3.1 Modeling process

To maximize crop yield and simultaneously optimize water and fertilizer use efficiency with limited environmental impacts, we needed to integrate a crop model with an optimization technique. The chosen crop model was the Decision Support System for Agrotechnology Transfer (DSSAT). DSSAT is a computer model capable of simulating crop growth for various cultivars and species. DSSAT considers the full cycle of soil-plant-atmosphere dynamics, irrigation scheduling, and nutrient management planning. The aforementioned characteristics along with the speed of the model (the whole growing season simulations taking few seconds) make this model ideal for this study. DSSAT can act as the evaluator within a bi-level evolutionary optimization framework.

In the bi-level optimization, one optimization problem is embedded (nested) in another. The outer and inner optimization problems are commonly referred to as upper- and lower-level optimization problems, respectively. Consequently, the variables of these problems are referred to as upper- and lower-level variables. The decision variables include two time-independent and two sets of time-dependent variables. The time-independent variables (x_i, $i \in \{1, 2\}$) represent the number of necessary irrigation and nutrient applications within a growing and are usually decided by an Irrigation Association and producers, respectively. The time-dependent variables ($y_j(t)$, $1 \leq j \leq \sum x_i, 1 < t < 365$) determine the amount of irrigation water and nutrient application for each date. We then evaluate the solution for a number of objectives: 1) maximizing crop yield, 2) minimizing irrigation water used, 3) minimizing the amount of nutrients applied, and 4) minimizing the nutrient loss (through leaching to shallow/groundwater and surface runoff).

Figure 2 illustrates the linking of the bi-level optimization, decision variables, and DSSAT crop model. The procedure starts with referencing the two time-independent/upper-level variables (number of irrigation and number of nutrient application events) in the bi-level optimization. A Monte Carlo optimizer uses these predefined variables to generate a series of uniformly distributed random day combinations within the growing season. For example, given three irrigation and two nutrient application events, one of the uniformly distributed random day combinations can be May 15th, July 18th, and August 1st for irrigation and May 30th and July 18th for nutrient applications. From the series of uniformly distributed random day combinations, one will be selected and incorporated into the lower level of the bi-level optimization. The lower level consists of the EMO algorithm Unified Non-dominated Sorting Genetic Algorithm-III (U-NSGA-III) (Seada and Deb, 2015). The algorithm U-NSGA-III can be used for single, multiple, and many-objective optimization problems alike, thereby allowing us to optimize systematically by starting with the most critical objective and then adding other objectives. The incorporated day combination will be used to initialize a population of time-dependent variables (irrigation and nutrient application amounts). These variables will be assigned to the dates identified in the upper level. For each solution in the population, the objective functions (O_1 to O_4) will evaluate these dates and amounts by calling the DSSAT model and retrieving the results. The population will be tested against termination criteria (e.g., predefined threshold values and maximum number of allowed iterations). If the termination criteria are not met, the population (the time-dependent variables) are evolved and re-evaluated. This procedure is repeated until the termination criteria are met, which then the local Pareto-front for the selected day combination will be stored. After each iteration of the upper-level of optimization, the new local Pareto is combined with a global Pareto front. In the next step, if there are any day combinations left, the above procedure would be repeated for each new day combination until all the generated random day combinations are processed.

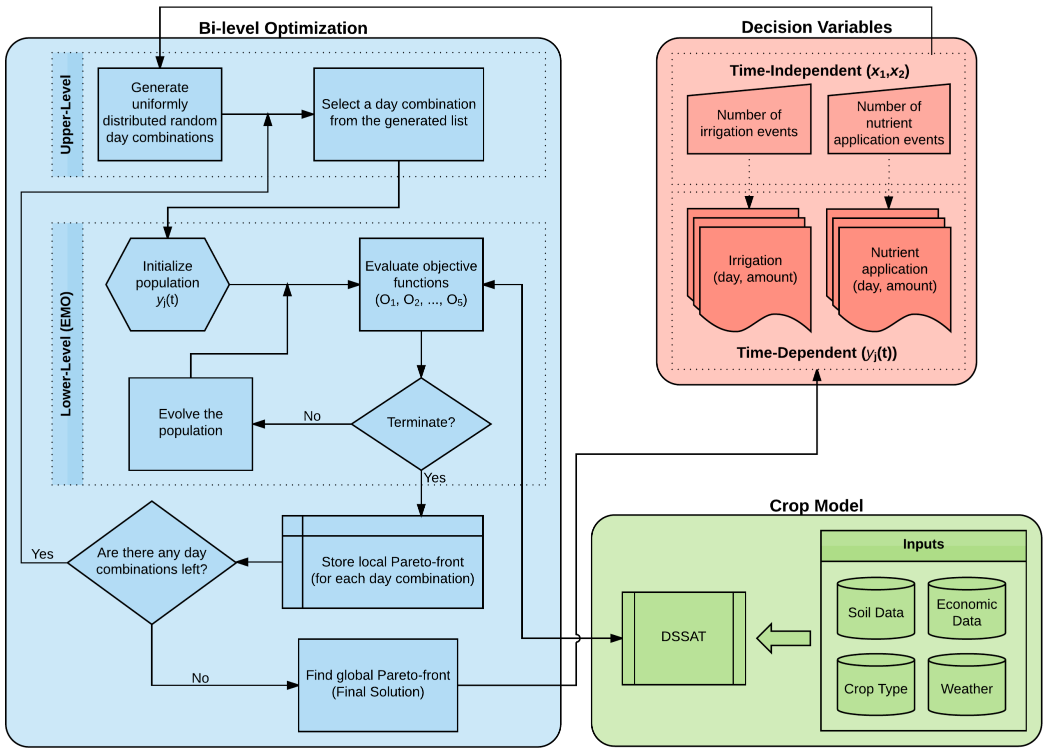

Figure 2. Proposed integrated bi-level and DSSAT framework

In order to compare the results from this study with an actual field experiment, the study consisted of two phases. In the first phase, we evaluated the efficacy of the lower-level optimization by examining the irrigation and nitrogen fertilizer application amounts for a fixed set of dates. The fixed set of dates followed the best management practice for high irrigation and high nitrogen fertilizer applications for a field experiment conducted at the Irrigation Research Park in Gainesville, Florida. Using the same application dates as best practices enabled accurate comparisons between the performance of the optimization and current best practices. In the second phase of the study, we tried to examine whether changing application dates and amounts from the fixed dates can further improve the overall performance of the farm considering the predefined objectives.

3.2 Study area

The study area is located at the University of Florida (UF) Experimental Station (29° 37.8', -82° 22.2') known as the Irrigation Research Park in Grasonville, Florida, USA (Figure 3). The study area is in a humid subtropical region of southeast of United States, with a wet season starting in May and ending in October (Lascody, 2002). The long-term (30 years) average, minimum, and maximum temperature of 20.7°C, 27.3 °C, and 14.6 °C, respectively. Meanwhile, the long-term total annual precipitation is 1,286 mm ("30-yr Normal Maximum Temperature: Annual," 2015, "30-yr Normal Mean Temperature: Annual," 2015, "30-yr Normal Minimum Temperature: Annual," 2015, "30-yr Normal Precipitation: Annual," 2015). In 1982 between February 16 and May 7*,* the mean daily maximum temperature was 26.96 C, the mean daily minimum temperature was 13.40 C, and the average daily rainfall was 5.07 mm. The dominant soil type of the field is Millhopper Fine Sand, which is moderately well drained. The experiment was conducted to evaluate the impacts of best management practices for irrigation and nitrogen fertilizer application on maximizing maize production (Hoogenboom et al., 2017). The study period occurs over one growing season during the spring and summer of 1982. Maize (McCurdy 84aa) was planted on February 16 and harvested at maturity on May 7. Among the six treatments in the experiment, we focused on treatment 4, which used irrigation and relatively high levels of nitrogen. Given the weather of 1982, the best practices from treatment 4 prescribed 16 days of furrow irrigation, totaling to 264 mm over the season. Nitrogen best practices from treatment 4 prescribed broadcasting ammonium nitrate incorporated at a depth of 10 cm depth six times over the growing season, totaling to 401 kg. This simulated experiment yields 11,298 kg/ha of Maize in DSSAT. The total seasonal yield, irrigation applied, and nitrogen applied, all act as benchmarks for current best practices during the optimization phase of this study.

Figure 3. Location of the field of study

3.3 Optimization Platform

3.3.1 Objective Function

Optimization problems alter one or more variables to maximize or minimize one or more problem objectives. In the case of crop production, producers alter agricultural variables (e.g., irrigation, fertilizer) to maximize their profits. This study focuses on the decisions of when to apply irrigation and/or nitrogen on a field, and how much irrigation and/or nitrogen to apply.

Summing the total seasonal irrigation and nitrogen applications provides simple assessments of environmental impact but estimating the expected seasonal yield is a more complex challenge. In fact, since researchers have dedicated years of research to developing software that can effectively model the performance (e.g. yield, nutrient leaching) of cereals and other agricultural staples (e.g. CROPGRO model (Boote et al., 1998), CERES-Maize (Jones et al., 1986), CERES-Wheat (Ritchie, 1985), SOYGRO (Wilkerson et al., 1983), PNUTGRO (Boote et al., 1992)) , instead of reinventing the wheel and redeveloping a mathematical model for crop yield, it is smarter to use an existing crop model as an objective function for yield in an optimization platform.

There are numerous existing crop models that could act as an optimization objective function. The DSSAT model is the clear best choice because it is fast, well validated, and easy to use (Chung et al., 2014; Hoogenboom et al., 2012). A user runs the model by defining the soil, weather, cultivar, and growing practices within a number of input files. The input files are then fed into the core of the DSSAT model, the crop simulation model (CSM). The CSM in return predicts a number of performance metrics for that hypothetical season (e.g., yield, nutrient usage, water use, nutrient leaching). The CSM itself consists of a highly modular system of sub-models that work together as a single unit. To name a few, DSSAT contains sub-models that handle the weather, soil (including nitrogen leaching), and numerous cultivars. And because researchers have validated the outputs of these sub-modules as a whole under numerous crop, climate, and soil conditions (Jones et al., 2003) Finally, DSSAT has been calibrated to simulate McCurdy 84aa maize growth at our study site in Gainesville Florida, and by extension is calibrated for the region and climate of Florida.

Using DSSAT, one can easily design a set of powerful crop production and/or environmental objective functions to use in an optimization algorithm, such as in our case:

$$[Max\colon Y, Min\colon L] = DSSAT\left(i_{\mathrm{A}0}, \ldots, i_{Aj}, i_{\mathrm{d}0}, \ldots, i_{dj}, f_{\mathrm{A}0}, \ldots, f_{Ak}, f_{\mathrm{d}0}, \ldots, f_{dk}\right) \quad (1)$$

$$Min\colon I = \sum_{n=0}^{j} i_{An} \quad (2)$$

$$Min\colon F = \sum_{m=0}^{k} f_{Am} \quad (3)$$

where, Y is yield, L is leaching, I is total irrigation, F is total fertilizer usage, i_{An} is the irrigation amount for date i_{dn}, f_{Am} is the nitrogen application for date f_{dm}, j is the total number of irrigation applications, and k is the total number of nitrogen applications. The optimization algorithms in this study contain combinations of these three objective functions.

All other variables (e.g., climate, soil, location), save irrigation application amount and nitrogen application amount, remained constant throughout the optimization. The units for irrigation applied was millimeter, the units for nitrogen applied was kilogram per hectare, and the units for nitrogen leaching is kg/ha. The total amounts of irrigation and nitrogen applied were assumed to be positive integers. Applications contained no upper bound as to allow the multi-objective optimization algorithm to find the full range of tradeoffs. Finally, all optimization runs used the same x_1 and x_2 values, or in other words, all strategies used the same number of irrigation and nitrogen applications. x_1 and x_2 both come from best field practices, in order to easily compare the performance of the optimization to best practices. However, often certain day values go to zero, which implicitly reduces x_1 and x_2 below best practices.

3.3.2 Optimization Algorithm Choice and Setup

This study optimizes crop management practices using multi-objective (MO) algorithms. MO algorithms are ideal for problems with conflicting objectives, such as minimizing irrigation and maximizing yield. The output of an n-objective MO algorithm is an n-dimensional Pareto set of solutions (Deb, 2009). These n-dimensional Pareto sets contain solutions that are all equally optimal and offer unique tradeoffs for each of the objectives. For example, one solution might maximize yield, but waste a lot of irrigation water. A second solution might minimize irrigation usage but yield very little product. These two solutions are both optimal with regards to separate objectives (irrigation or yield) and are equally optimal. A decision maker, such as a farmer or an environmental policy maker, would then choose the solution that best fits their need.

Building on the concepts of MO algorithms, EMO algorithms incorporate the principles of evolutionary optimization with MO concepts. Evolutionary optimization algorithms solve problems by creating populations of solutions (or individuals), ranking the individuals, and then recombining the best solutions into a new and ideally improved generation of solutions.

Evolutionary algorithms excel at solving complex, nonlinear objectives (Sastry et al., 2014), and are therefore a powerful tool for plant and crop optimization.

The EMO algorithm NSGA-II efficiently sorts individuals by dominance and then by diversity (Deb et al., 2002). Given enough generations, it yields maximum convergence (closeness of the Pareto front to one or more objectives) and diversity (the spread of the Pareto front) for the given problem and seed. NSGA-II's prioritization of both performance metrics cuts it out among similar EMO algorithms. Most decision makers seek the most optimal solutions and equally seek diverse choices in their decision-making process. However, NSGA-II performs optimally between one to three objectives. Meanwhile, crop production usually requires meeting many objectives, so this study utilized the U-NSGA-III algorithm. U-NSGA-III works similarly to NSGA-II, but it improves the crowding distance calculations of NSGA-II to effectively calculate crowding in one- to many-dimensional objective spaces (Deb and Jain, 2014); (Seada and Deb, 2015).

U-NSGA-III was set to optimize with real variables; however, because the variables in this agricultural optimization system are all integers, the variables were normalized between 0 and 1. Zero represented zero, and 1 represented the largest value of that variable type. Normalizing between these two relatively smaller values reduced the variable search space for U-NSGA-III. For example, for irrigation, 0 represented 0 mm of irrigation applied, and 1 represented 50 mm of irrigation applied. For nitrogen, 0 represented 0 kg/ha, and 1 represented 150 kg/ha. Software within the objective function scales these variables back up and rounds them to the nearest whole number for DSSAT.

The COIN Laboratory at Michigan State University provided the Java U-NSGA-III implementation, named *evolib* (Seada, 2017). At the commencement of this research project, *evolib* was single threaded and to make use of the 24 core processor in our lab computer; we forked *evolib* to include multithreaded evaluation of population (Seada and Kropp, 2018). Additional software development efforts included the project DSSAT4j (Kropp, 2018a). Since the DSSAT CSM is written in FORTRAN, much of the project focused on writing a Java wrapper for DSSAT CSM input and output. The wrapper provides abstracted DSSAT experiments in the form of Java objects. For example, a client of the DSSAT4j library might create an experiment object containing an array of dates and irrigation amounts. A simple call to a run method returns the yield and/or leaching for the given experiment. This allowed simple and elegant integration of DSSAT into a U-NSGA-III objective function. DSSAT4j also includes built-in multithreading, which enabled quick, multi-threaded U-NSGA-III. DSSAT4j ran DSSAT version 4.7.1 (Git commit d8a977d) (Hoogenboom et al., 2017). The final component of the software framework is software that tied it all together. This software, named CropOpt (Kropp, 2018b), is a prototype decision support framework and optimizes a crop growing scenario provided by a user. The user provides periods of irrigation and nitrogen application, along with constraints such as the number of irrigation and nitrogen application, maximum and minimum amounts applied per application, and so on. In return, the framework optimizes the scenario and returns the results to the user. All results can be found in the supplementary data in the DSI website (see data availability).

3.3.3 Optimization Strategies and Configuration

The study chose to use a bi-level optimization configuration over single-level optimization because of the complexity of the problem. The number of variables behind irrigation and nutrient scheduling are numerous and difficult to optimize. If a hypothetical growing season contains 120 days, and a farmer has access to 0 to 50 mm of water every day, there are 51^{120} or $8.09 \cdot 10^{204}$, different solutions. In its early stages, this study considered every day between February 16 and harvested the maize at maturity on May 7 a real variable in the optimization problem. However, the optimization runs with this single level optimization configuration failed to match the UF field experiment benchmark yield. Therefore, the study sought to reduce the number of variables while still optimizing for the 120-day growing season with a bi-level optimization scheme. The upper level of the optimization was a Monte Carlo optimization algorithm, which generated random combinations of dates to irrigate and/or to apply nitrogen. For each of these combinations, U-NSGA-III attempted to find the optimal application amounts for each fixed date combination. Formally, the algorithm ran as follows:

Algorithm 1: Bi-Level Optimization for Agricultural Optimization

Step 1: Determine the number of irrigation applications x_1

Step 2: Determine the number of nitrogen applications x_2

Step 3: Repeat

Step 4: Generate x_1 random dates for irrigation application

Step 5: Generate x_2 random dates for nitrogen application

Step 6: Run U-NSGA-III with the application amounts of given irrigation and nitrogen dates as variables

Step 7: Combine the solutions from step 6 with solutions found in previous iterations

Step 8: Remove dominated solutions from the combined set of solutions

Here we hypothesized that through bi-level optimization, better management practices could be implemented to not only improve crop yield, but to also minimize irrigation application rates, minimize fertilizer application rates, and minimize environmental impacts through leaching. To verify this hypothesis, different optimization strategies were formulated and evaluated based on their ability to improve current best management practices. As a result, each optimization strategy maximized yield while minimizing a different combination of environmental impacts, including total irrigation applied, total nitrogen applied, and total nitrogen leached. The first three optimization strategies use single-level optimization with dates taken from UF best experimental field practices, and the last three use bi-level optimization.

Table 1. Summary of the optimization strategies

Strategy No.	Optimization Type	Application Dates	Objectives
1	**Single Level** (U-NSGA-III)	16 fixed irrigation dates (dates from the UF field experiment)	Max: Yield Min: Irrigation applied
2	**Single Level** (U-NSGA-III)	16 fixed irrigation dates AND 6 fixed nitrogen dates (all dates from the UF field experiment)	Max: Yield Min: Irrigation applied Min: Nitrogen applied
3	**Single Level** (U-NSGA-III)	16 fixed irrigation dates AND 6 fixed nitrogen dates (all dates from the UF field experiment)	Max: Yield Min: Irrigation applied Min: Nitrogen applied Min: Nitrogen leached
4	**Bi-Level** (Monte Carlo for upper-lever & U-NSGA-III for lower-level)	16 variable irrigation dates	Max: Yield Min: Irrigation applied
5	**Bi-Level** (Monte Carlo for upper-lever & U-NSGA-III for lower-level)	16 variable irrigation dates 6 variable nitrogen dates	Max: Yield; Min: Irrigation Min: Nitrogen applied
6	**Bi-Level** (Monte Carlo for upper-lever & U-NSGA-III for lower-level)	16 variable irrigation dates 6 variable nitrogen dates	Max: Yield; Min: Irrigation Min: Nitrogen applied Min: Nitrogen leached

The various strategies are as follows. *Strategy 1* kept all parameters from the UF experiment constant, including the 16 irrigation dates, but changes the irrigation application rate for each date. Keeping all other parameters constant made the results easily comparable to current best practices. *Strategy 2* similarly kept all parameters constant except for irrigation and nitrogen amounts. Like *Strategy 1, Strategy 2* uses the same dates of irrigation application and dates of nitrogen application as best practices. *Strategy 3* was a four-dimensional optimization problem, and contained the same variables as *Strategy 2*, but added an additional objective to minimize nitrogen leaching. *Strategy*

4, the first bi-level optimization run, tried to find both the dates for 16 irrigation applications and the amounts for each of those irrigation applications. The objectives of *Strategy 4* mirror the objectives of *Strategy 1. Strategy 5,* which mirrors the objectives of *Strategy 2,* uses bi-level optimization to find the dates and amounts for 16 days of irrigation application and to find the dates and amounts for six days of nitrogen application. The final bi-level optimization strategy, *Strategy 6*, which mirrors the objectives from *Strategy 3,* has the same variables as *Strategy 5* while adding a leaching minimization objective to the problem.

There are two main stopping criteria employed in the study. In the first stopping criteria, *Strategies 1* through *3* ran until their hypervolumes visually plateaued. A hypervolume is the volume of the Pareto front from a reference point (Emmerich et al., 2005). The reference point is typically a worst-case scenario value for each objective, and in this study, we chose the following worst-case values in our reference points: 1) Yield: 0 kg/ha 2) Irrigation: 1000 mm 3) Nitrogen: 1000 kg/ha 4) Leaching: 500 kg/ha. Once the hypervolume plateaued, the value with the greatest hypervolume was chosen as the best solution. In the second stopping criteria, *strategies 4* through *6* ran U-NSGA-III until a fixed generation. The fixed generation is a predetermined generation at which the hypervolumes of *Strategies 1* through *3* respectively stopped improving. The study used the hypervolume software developed by Walking Fish Group, which can calculate 1 to many dimension hypervolumes (Cox and While, 2016).

3.4 Post-processing and visualizations

Once validated, the study analyzes the solutions within the available Pareto fronts. The first step of analysis was to break the Pareto fronts into three clusters. Three clusters effectively break the data into three broad categories: high productivity, high environmental efficiency, and even trade-off. All tradeoffs in every optimization scenario can be broken down in these categories because the tradeoffs are all essential bipolar. Yield maximization falls under the productivity pole, and irrigation minimization, nitrogen minimization, and leaching minimization all fall under the environmental efficiency pole. The third category attempts to find solutions that are well balanced between these poles. This study employs the *k*-means clustering method to break down the fronts into three categories (Arthur and Vassilvitskii, 2007). For this study, the *k*-means algorithm separated the solutions based on their normalized objective values. However, other studies could weigh the objectives differently according to the needs of a decision maker.

The study also developed a unique type of histogram to display the clustered data (section 3.3). Each bucket in the histogram contains stacked, color-coded bars for each of the clusters. This simultaneously represents the total distribution of the Pareto front as well as the distribution of the individual clusters in the same front. The application frequency analysis section uses this type of histogram.

Regarding the visualization, *Strategy 6* posed a challenge. A three-dimensional graph can effectively display the results of *Strategies* 1 through 5 because each strategy contains 2 to 3 objectives. However, *Strategy 6*, on the other hand, attempted to optimize four objectives, which cannot exist as points on a three-dimensional space. To address this problem, a custom written MATLAB script first plots the first three dimensions of the four-dimensional solutions. Then, the script color-codes the three-dimensional points with the values of the fourth dimension. This method effectively conveys the four-dimensional objective space in a three-dimensional objective space.

4. RESULTS AND DISCUSSIONS

4.1 Single-level optimization results

4.1.1 Strategy 1: Single Level Irrigation Minimization and Yield Maximization

Strategy 1 optimized the simplest set of objectives. It attempted to maximize yield and minimize irrigation applied. The optimization found a diverse set of solutions that required significantly less irrigation than the UF best practices. The solution with the highest yield matched the output of the best practices (11,298 kg/ha) while reducing the required irrigation amount by 45.5% (Table 2). However, this solution is arguably not properly Pareto efficient, because the tradeoff of irrigation to yield is heavy. Properly Pareto efficient solutions are solutions that are Pareto optimal and offer reasonable tradeoff compared to their peers (Geoffrion, 1968). For example, choosing between solution A (Figure 4), which yields 11,298 kg/ha with 136 mm of irrigation, and solution B, which yields 11,242kg/ha with 130 mm of irrigation, is intuitively simple. Moving from solution B to solution A gains just 0.49% while increasing total irrigation by 4.4%. Therefore, solution A is arguably not properly Pareto efficient for many stakeholders. However, the final call of what is properly Pareto optimal or not belongs to the stakeholders themselves, and to a farmer who has no water use limit, both are arguably properly Pareto efficient. Meanwhile, the solution with the lowest yield produced 2,655 kg/ha with 0 mm of irrigation.

The optimization also significantly reduced the number of irrigation applications. In fact, all solutions required 13 or fewer applications during the season. The highest yielding solution required 11 applications, which is a noticeable improvement over the best practices that required 16 applications to achieve the same yield (11,298 kg/ha). Farmers with regional restricted temporal water access would benefit using solutions that reduce the number of applications while maintaining optimal yield (Sampath, 1992). This is an example of how the optimization can shape management practices. The number of applications within the Pareto front (ranked in descending order by yield) non-monotonically decrease from 11 down to zero.

Figure 4 illustrates this transition between high yielding solutions and highly water efficient solutions. The more water is conserved between solutions, the more yield decreases. Moving from low water solutions to high water solutions at first offers relatively high yield tradeoffs for low water tradeoffs. However, approximately halfway through the solutions, unit increases in total irrigation cease to offer equal increases in total yield (Figure 4).

Table 2. Top five (ranked by yield) optimal results compared to best practices (*Strategy 1*)

Yield (kg/ha)		Number of Applications		Total Irrigation (mm)		
Best practices	Optimized results	Best practices	EMO results	Best practices	EMO results	Reduction in irrigation usage
11298	11298	16	11	264	136	48.5%
	11267		11		130	50.8%
	11242		11		129	51.1%
	11204		11		125	52.7%
	11128		13		119	54.9%
	11075		11		115	56.4%

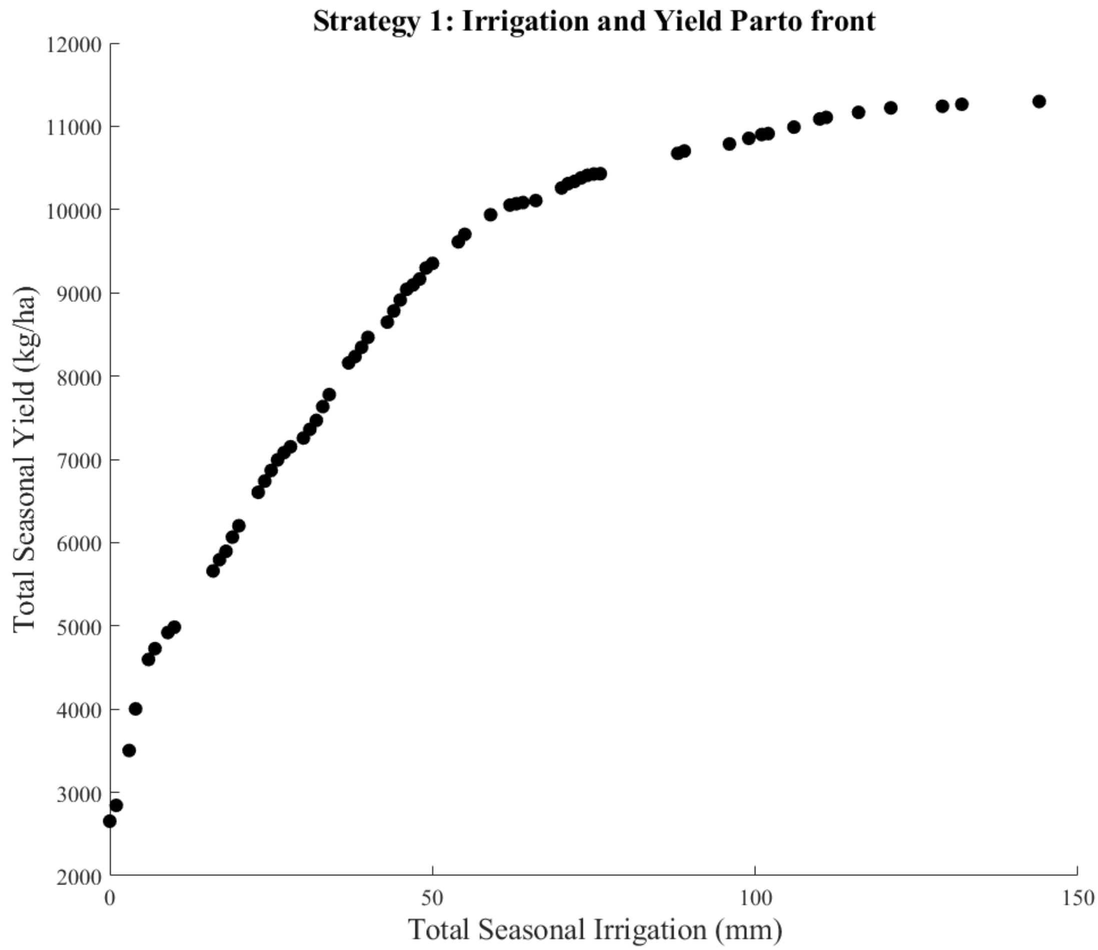

Figure 4. Pareto front for *Strategy 1*, which maximized yield and minimized irrigation amount

4.1.2 Strategy 2: Single Level Irrigation Minimization, Nitrogen Minimization, and Yield Maximization

Strategy 2 added an objective to minimize nitrogen usage to the objectives of *Strategy 1* and obtained similar reductions in nitrogen and irrigation amounts (Table 3). The best solution with respect to yield (top row of Table 3) also matched the yield benchmark from the UF best practices, while reducing irrigation by 40.9% and nitrogen by 26.4%. The lowest ranked solution by yield applied 3 mm of irrigation and 7 kg/ha of nitrogen and reached 939 kg/ha of yield. This worse case (by yield) solution from *Strategy 2* performs considerably worse with respect to yield than the worst case from *Strategy 1* because *Strategy 1* used the nitrogen application dates and amounts from the best practices, and therefore does not use more or less than 401 kg/ha. Though the yield is lower than best practices, this yield is still optimal given harsh irrigation and fertilizer restrictions because it is Pareto optimal (Deb, 2009).

Despite their relatively low yields, examining the lowest solutions by irrigation and nitrogen still provides valuable agricultural insights. The lowest ranked solution by nitrogen applied 3 kg/ha of nitrogen and 74 mm of irrigation to achieve 1566 kg/ha of yield. The two lowest ranked solutions by irrigation both apply 3 mm of irrigation, apply 7 kg/ha and 41 kg/ha of nitrogen respectively, and achieved 989 kg/ha and 1778 kg/ha in yield respectively. Both of these cases offer management practices to farmers under severe but common scenarios. A hypothetical farmer who wishes to sell organic maize may choose the solution that uses almost no nitrogen, and a farmer in an area with strict irrigation restrictions may choose a solution that uses almost no irrigation. In both cases, the Pareto front offers the optimal yield found in the optimization under such restrictions.

The solutions that contain the highest irrigation and nitrogen usage also offer valuable tradeoffs. The highest ranked solution by nitrogen actually applied more total nitrogen (500 kg/ha) than the best practices (401 kg/ha) to produce 8,683 kg/ha of yield. However, the solution makes up for its heavy nitrogen usage by reducing the irrigation applied. A similarly yielding solution in *Strategy 2* (8,593 kg/ha) uses 68.9% more irrigation while using less nitrogen (123 kg/ha). Therefore, a stakeholder may choose to save a great deal of irrigation by choosing the management practice that uses a relatively large amount of nitrogen fertilizer. The highest solutions ranked by irrigation, using 385 mm of irrigation and 186 kg/ha of nitrogen to yield 10,870 kg/ha, offers a similarly unique trade-off. Compared to a similar solution in *Strategy 1* that yields 10,859 kg/ha with 401 kg/ha of nitrogen and 101 mm of irrigation, this solution from *Strategy 2* decreases nitrogen usage by 56.61% while sacrificing considerable quantities of water. These cases offer unique solutions for unique scenarios and give stakeholders the final judgment whether to sacrifice one objective to improve two or more other objectives. These results also demonstrate the broadening of options in optimizing both irrigation and nitrogen application within the same optimization run. With the third objective of nitrogen, stakeholders not only determine the appropriate amount of water but also the appropriate amount of nitrogen as well.

The above tradeoffs are visually present in the *Strategy 2* Pareto front in Figure 5. The top of the front, where water and nitrogen are highest, contains the highest yields, but as you conserve more nitrogen and water, the yield decreases at the skirts of the front.

The number of nitrogen applications in *Strategy 2* does not decrease as dramatically as the irrigation applications in *Strategy 1.* When the data is broken into high, medium and low yield clusters (using *k*-means clustering), all the high yield solutions were within 4 to 6 applications. Furthermore, the solutions using 4 to 5 applications were outside of the interquartile range of the high yield cluster. Even for the medium and low clusters, all solutions between the first quartile and the maximum were between 3 and 6 applications (Figure A.1 in the Supplementary Materials). This suggests that solutions using under three nitrogen application have a higher risk of being dominated by other solutions. Also, the fact that the median, the third quartile, and the maximum are all the same value in the high yield cluster (Figure A.1) suggest that solutions with more than six nitrogen applications may contain competitive contributions to the existing Pareto front. However, this paper focuses on the optimization problem using a maximum of six possible days to keep in line with the UF benchmarks field experiment and seeks to keep the nitrogen application count within realistic bounds. Until the economic impacts of increasing the number of nitrogen applications above 6 can be taken into account in future studies, we assume that the maximum number of nitrogen application is six.

The number of irrigation applications does not appear to directly affect yield as strongly as the number of nitrogen applications (Figure A.1). In fact, in the high yield cluster, the range is very wide (4 applications to 16 applications), though the middle 50% of the solutions fall between the narrower range of 9 to 14 applications. The medium yield cluster has slightly a wider range (3 applications to 16 applications) and a higher middle 50% range (13 applications to 5 applications), and the low yield cluster has a similar range and inner quartile range (2 to 15 applications and 5 to 12 applications respectively). These increasing ranges suggest that the effects of the number of irrigation applications have a diminishing effect on the yield as the target yield decreases.

Similar to *Strategy 1,* the number of irrigation applications on the highest-ranking solution by yield uses fewer irrigation applications than best practices (14 instead of 16). However, 21 solutions still required 12 to 16 applications. The added complexity of six additional variables and one additional objective from *Strategy 1* might explain the presence of these solutions with high application counts. Adding more variables increases the dimensionality of the decision space and adding more objectives increases the dimensionality of the objective space. Both cases increase the difficulty of finding optimal solutions (Deb, 2009). Alternatively, perhaps more irrigation applications are optimal when simultaneously optimizing nitrogen and irrigation.

Another interesting finding in Figure 5 is that there is an upward trend in yield from low irrigation and nitrogen application to high irrigation and nitrogen application. The trend dictates that with every horizontal cut, along the yield axis, there is one trade-off solution that can be used to help farmers to make better decisions. For example, if we extract the green shade from the figure, we have a two-dimensional Pareto front with respects to irrigation and nitrogen. Along this Pareto front, there is a point known as the knee point, which offers a trade-off solution among irrigation and nitrogen application. Similar points can be created at different yield levels and can be presented to farmers to simplify the decision-making process.

Table 3. Top five (ranked by yield) optimal irrigation and nitrogen results compared to best practices (*Strategy* 2).

Yield (kg/ha)		Total Irrigation (mm)		Irrigation Reduced	Total (kg/ha)	Nitrogen	Nitrogen Reduced
Best practices	EMO results	Best practices	EMO results		Best practices	EMO results	
11298	11298	264	156	40.9%	401	295	26.4%
	11208		223	15.5%		219	45.4%
	11142		127	51.9%		325	19%
	11106		291	-10.2%		203	49.4%
	11102		164	37.9%		281	29.9%

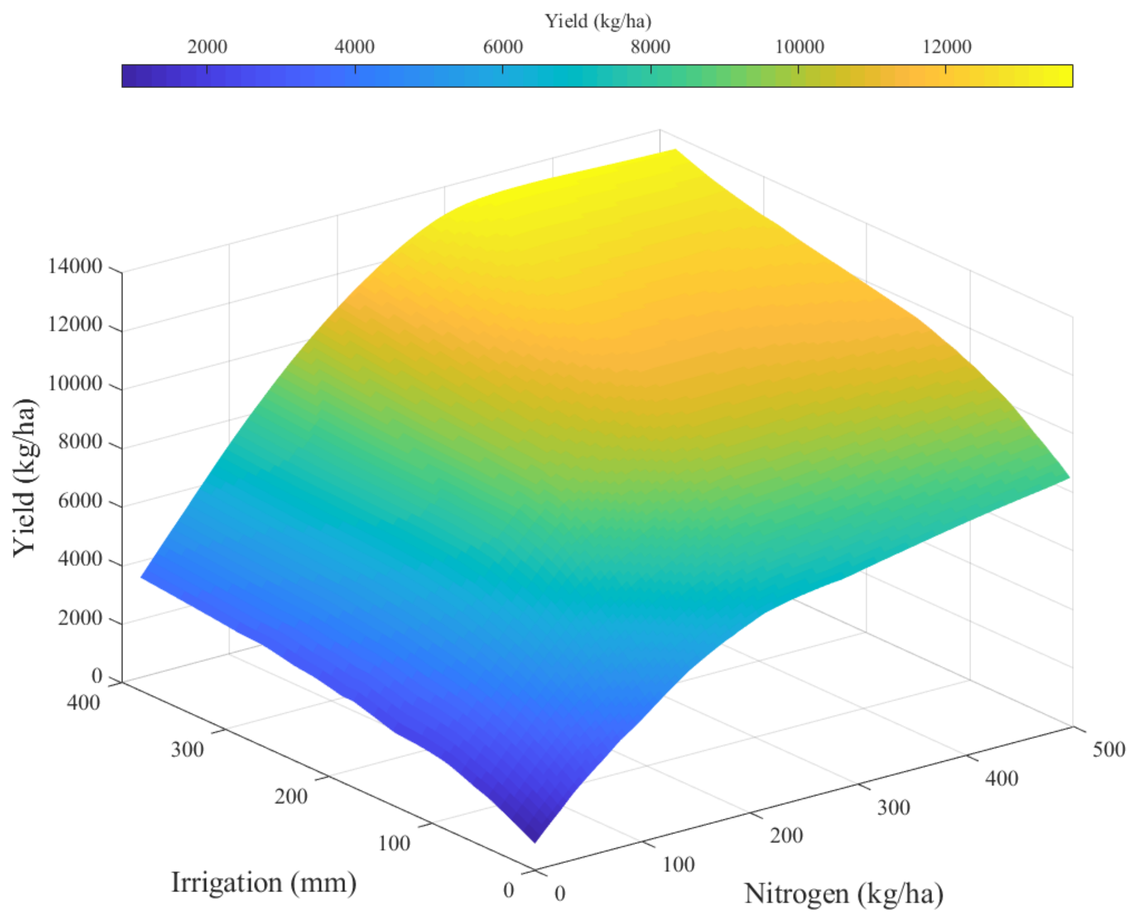

Figure 5. Pareto front for Strategy 2. A surface was passed through the solutions to better visualize the shape of the front. The color of the surface represents the yield for that given region of the front.

4.1.3 Strategy 3: Single Level Irrigation Minimization, Nitrogen Minimization, Leaching Minimization and Yield Maximization

Strategy 3 includes the effects of a direct environmental impact: nitrogen leaching. Nitrogen leaching occurs when nitrogen-based fertilizers infiltrate the soil and ultimately the groundwater.

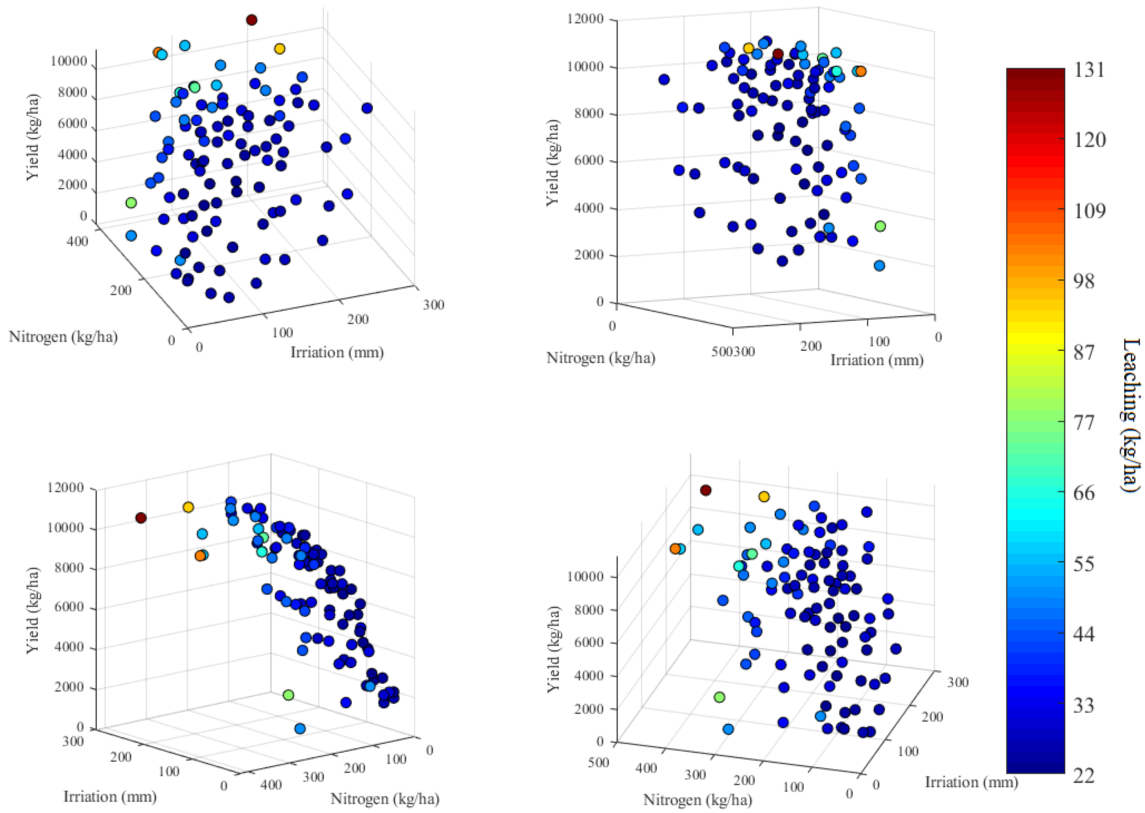

Figure 6. Pareto front for *Strategy 3,* which maximized yield, minimized nitrogen application, minimized irrigation application, and minimized nitrogen leaching. The color of the surface represents the leaching of the solution

To incorporate this environmental metric, *Strategy 3* included four objectives: minimize nitrogen application, minimize irrigation application, minimize nitrogen leaching, and maximize yield.

The optimization results accurately reflect the literature on leaching. The results with relatively high irrigation application and maximum nitrogen applications yielded the most maize, but also the most leaching. However, the results provide solutions that reduce dangerous leaching levels by simply reducing the yield by approximately 100 kg/ha. Knowing this information is important for decision-makers to incentivize farmers to avoid over-fertilizing their fields, and thus reducing leached nitrogen in water supplies.

The yield, water, and nutrient usage among high yielding solutions vary only slightly from *Strategy 1* and *Strategy 2.* The best Pareto front of the run contained solutions yielding up to 11,281 kg/ha (17 kg/ha lower than the highest yielding solutions from *Strategy 1* and *Strategy 2*). Also, the same solution that achieved 11281 kg/ha used more considerably more irrigation than the highest yielding solutions in *Strategy 1* and *Strategy 2*. Specifically, *Strategy 1* and *Strategy 2* both yielded 11,298 kg/ha with 134 mm of irrigation and 156 mm respectively, where *Strategy 3* needed 180 mm of irrigation to achieve 11,281 kg/ha. The amount irrigation appears to trend upwards with each additional objective added, with the *Strategy 3* solution using the most. However, neither the *Strategy 1* solution nor the *Strategy 2* solution dominates this *Strategy 3* solution because the *Strategy 3* solution outperforms both of them in total nitrogen applied and leaching. *Strategy 3* achieves 11,281 kg/ha, approximately the same yield as *Strategy 1* and *Strategy 2*, but with only 285 kg/ha of nitrogen applied and 65. 6 kg of nitrogen leached. The corresponding solutions from *Strategy 1* and *Strategy 2* applied 401 kg/ha (from best practices) and 295 kg/ha of nitrogen respectively and leached 75.2 kg/ha and 92.6 kg/ha respectively. This demonstrates the utility in adding additional objectives to the optimization. The more objectives included in our optimization platform, the nuanced the possible management practices become.

The remaining value ranges were in most ways similar to *Strategies 1* and *2*. Yield ranged from 1,252 to 11,281 kg/ha. Leaching ranged from 21.99 to 119.5 kg/ha. Irrigation applications ranged from 8 mm to 276 mm, and nitrogen applications ranged from 9 kg/ha to 558 kg/ha. The number of irrigation applications ranged from 2 to 16 applications, but unlike *Strategy 1* and *2,* the number of nitrogen applications ranged from 4 applications to 6 applications. This suggests that when treating leaching as an objective, solutions with 0 to 3 nitrogen applications are dominated by their 4 to 6 application counterparts.

4.2 Bi-level optimization results

4.2.1 Strategy 4: Bi-Level Irrigation Minimization and Yield Maximization

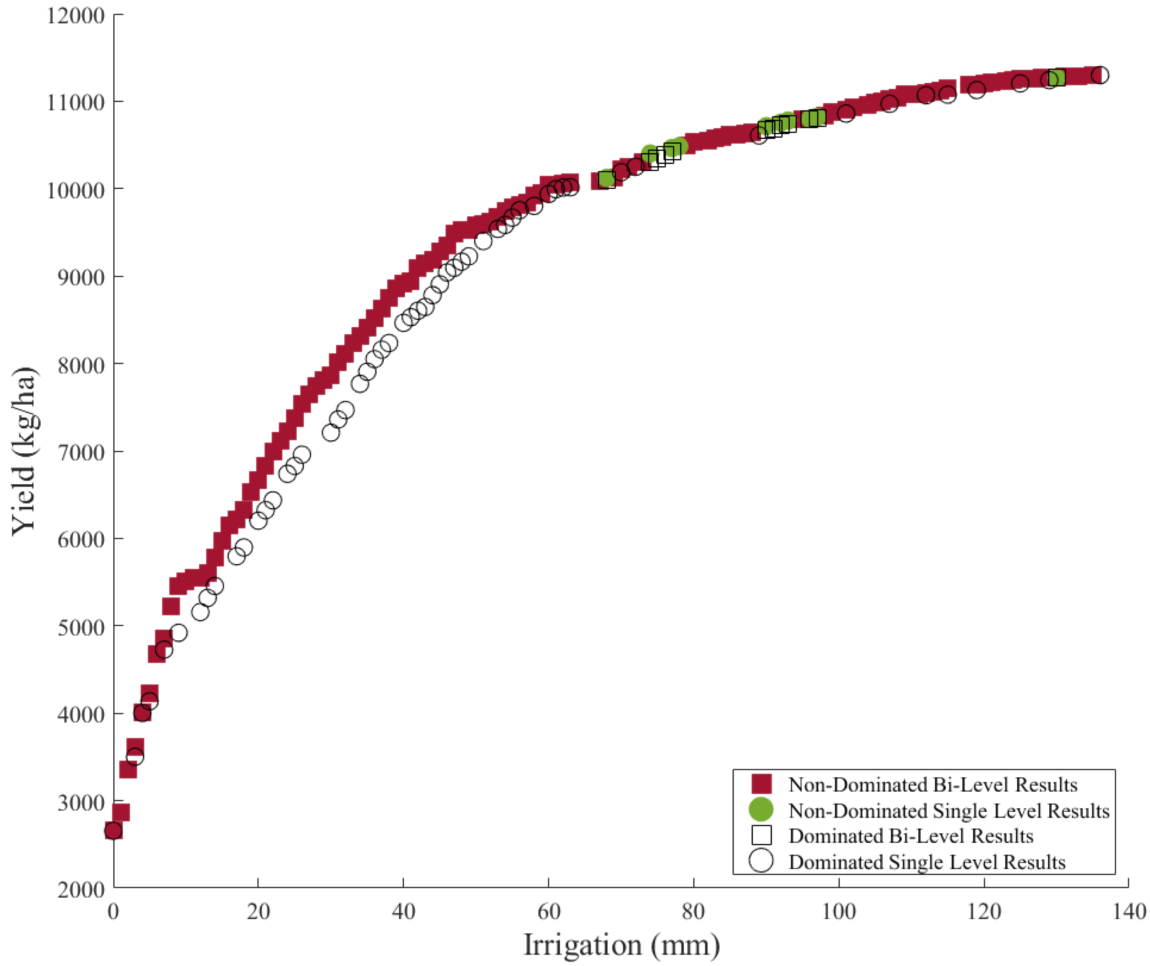

Figure 7. *Strategy 1* results compared to *Strategy 4* results

The combined Pareto front of the 372 separate optimization runs, consisting of 119 Pareto optimal solutions, outperformed the optimization from *Strategy 1*. Firstly, the hypervolume (using 0 kg/ha of yield and 1,000 mm of irrigation as a reference point) for the bi-level optimization was 11,027,186, where the hypervolume for the single level was 10,998,189. The larger hypervolume suggests better overall convergence and diversity from the bi-level results over the single level results (Deb, 2009). This improvement is visually present in Figure 7. Secondly, the middle segment of the *Strategy 4* Pareto front is significantly more convergent than the middle of the *Strategy 1* Pareto front. More convergence in a Pareto front means more points are closer to the ideal of maximizing irrigation and minimizing irrigation (Figure 7) (Rudolph, 1994). However, it worth noting that *Strategy 1* solutions dominate 10% of *Strategy 4* solutions, and *Strategy 4* solutions dominate 88% of *Strategy 1* solutions. But as Figure 7 illustrates, the dominated *Strategy 4* solutions only marginally derivate from their *Strategy 1* neighbors.

There are possible explanations to why *Strategy 1* still dominates a few *Strategy 4* solutions, and why *Strategy 4* solutions dominate the vast majority of *Strategy 1* solutions. The overall trend is that the single level run from *Strategy 1* marginally dominates in high yielding areas of the front, where *Strategy 4* solutions dominate the most dramatically in the middle of the front. *Strategy 1* holds out in high yielding areas of the Pareto front because *Strategy 1* already has dates chosen by experts to maximize yield. Therefore, it excels in finding solutions in this reduced search space. On the other hand, *Strategy 4* likely outperforms *Strategy 1* in the middle of the front because it opens the search space to date combinations that better suit median yield values. Both *Strategy 4* and *Strategy 1* performed similarly in the lower yielding solutions at the bottom of the Pareto because minimizing yield simply involves reducing irrigation levels to zero for every date. These runs clearly demonstrate the usefulness of this bi-level optimization strategy. The Monte Carlo optimization layer widens the narrow manual optimization of dates beyond yield maximization and provides stakeholders with more refined management practices to their particular cases, from low to high yield, and low to high irrigation availability.

4.2.2 Strategy 5: Bi-Level Irrigation Minimization, Nitrogen Minimization, and Yield Maximization

Like *Strategy 4*, *Strategy* 5 found a better Pareto front than the single level optimization Pareto front. Firstly, the number of solutions in the front increased dramatically. Maintaining a running bank of non-dominated points after each generation and after each run resulted in a 2,971 solution Pareto front after 137 runs. This differs significantly from the 119 solutions found in *Strategy 4* and is due to the additional dimension in objective space (nitrogen application amount minimization). Offering more optimal solutions to decision makers allows them to make decisions that are more refined after optimization. Secondly, the solutions from *Strategy 5* dominated the vast majority of single-level irrigation, nitrogen and yield optimization solutions from *Strategy 2.* In fact, 98% of solutions from *Strategy 2* are dominated by *Strategy 5*, and only 0.2% of *Strategy 5* solutions were dominated by *Strategy 2* solutions. Finally, the hypervolume of the single level optimization, with a reference point at 0 kg/ha of yield, 1,000 mm of irrigation, and 1000 kg/ha of nitrogen, is 1.01· 1010. The hypervolume of the bi-level optimization, using the same reference points, is slightly larger at 1.03· 1010, which suggests better over convergence and/or diversity of the front from *Strategy 5* compared to the front from *Strategy 2.*

An interesting difference between the two objective bi-level optimization run (*Strategy 4)* and the three objective bi-level optimization run (*Strategy 5)* is in how they respectively outperformed the single level optimization runs. The single level solutions from *Strategy 1* that outperformed the bi-level results from *Strategy 4* were mainly focused on in the high yield region of the Pareto front. However, when an additional objective is added in *Strategy 5,* the only single level solutions from *Strategy 2* that outperforms bi-level results from *Strategy 5* are two seemingly random spots on the Pareto front. Another major difference between the performance of *Strategy 4* and *Strategy 5* is the number of runs necessary to overcome their single level strategies. *Strategy 5* was able to dominate a greater amount of *Strategy 2* with fewer runs (137 runs) than *Strategy 4* did in 372 runs. One explanation for these differences is that since *Strategy 2* had one more objective than *Strategy 1* and thus had a more complex objective space to traverse. Therefore, *Strategy 2* struggled to converge on the potential optimum within a considerably large search space, and could not compete with the solutions from the bi-level optimization from *Strategy 5*. This demonstrates the power of the bi-level Monte Carlo and U-NSGA-III optimization that aids in avoiding local optima during optimization and increases the breadth of the Pareto front of management practices at the end of optimization.

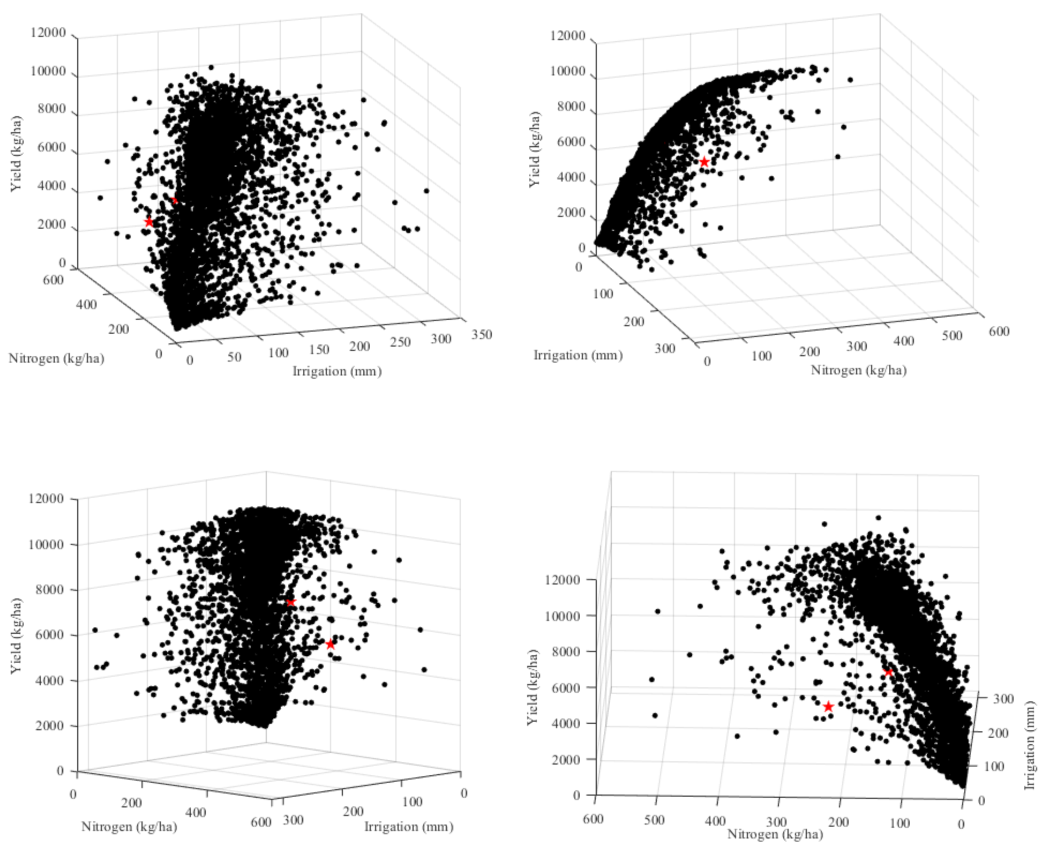

Figure 8. *Strategy 2* versus *Strategy 5* results. Non-dominated *Strategy 5* results are the black circles and non-dominated *Strategy 2* results are red stars

4.2.3 Strategy 6: Bi-Level Single Level Irrigation Minimization, Nitrogen Minimization, Leaching Minimization and Yield Maximization

Strategy 6 continued the trend of *Strategy 4* and *5*: it took fewer runs (101) to dominate even more of its corresponding single level run (*Strategy 3*). *Strategy 6* dominated all but one the solutions from *Strategy 3,* and the single non-dominated *Strategy 3* solution only dominated 0.2% of the *Strategy 6* solutions. Again, the single level optimization version of this four-objective problem (*Strategy 3)* was ill-equipped to converge on the full front found by *Strategy 6.* The addition of the fourth objective and the 6 additional variables for leaching (compared to *Strategy 1* and *4*) is the source of the additional complexity. One final benefit of *Strategy 6 is* even a greater number of solutions (6,198) found caused by the additional fourth dimension of objective space. The large number of solutions better defines the shape of the Pareto front than *Strategy 3*. *Strategy 3* lacks a well-defined shape, but *Strategy 6* reveals that the tradeoffs of the four-objective optimization problem are very similar to the tradeoffs of the three-objective problems in *Strategy 2* and *Strategy 5.* However, the shape of the *Strategy 6* Pareto front adds additional insights into environmental agricultural optimization. No solutions exist at the top of the front, where irrigation approximately exceeds 150 mm and nitrogen application roughly exceeds 250 kg/ha. The very sharp, near 90-degree-shaped boundary suggests that past a certain amount of combined irrigation and nitrogen application, no optimal solutions exist. This piece of knowledge could aid farmers and policymakers in making decisions that avoid both adversely affecting the environment and adversely affecting the yield of the crop.

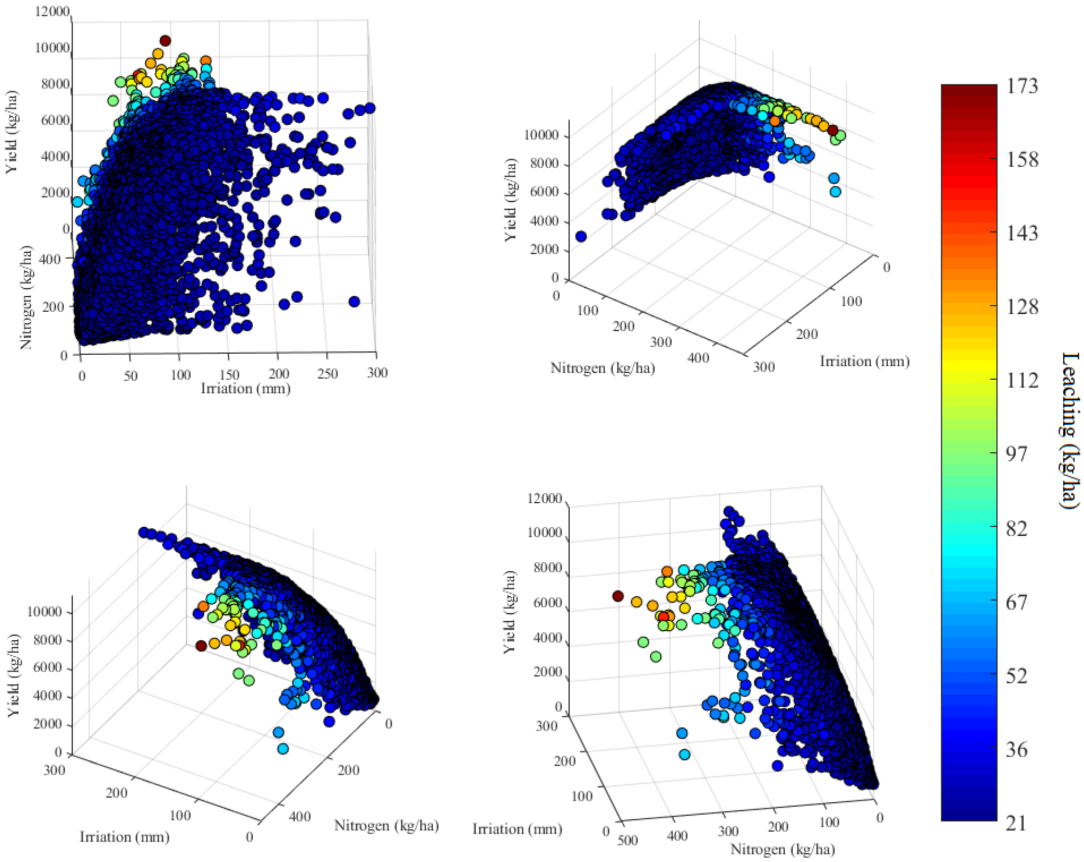

Figure 9. Non-dominated *Strategy 6* results

4.3 Application frequency analysis

Once *Strategies 4* through *6* found Pareto optimal fronts for their respective scenarios, the study categorized and analyzed trends within the results of each strategy. As discussed in section 2.4, the *k*-means clustering method partitioned the results into three regions: high yielding solutions, environmentally efficient solutions, and even trade-off solutions. The distribution of application counts required within each cluster is a powerful tool for developing management practices. For example, if a farmer only has access to five applications of irrigation in a season, he/she can estimate the expected yield in the season and how he/she can implement the practice. It can also help reduce the complexity of future optimization problems. For example, if a farmer seeks to achieve a yield within the medium trade-off cluster, and still wants to optimize irrigation usage, a new optimization routine would only need 6 irrigation variables instead of 16. Reducing an optimization problem from 16 to 6 would significantly reduce the complexity and run time of an optimization routine (Deb, 2009).

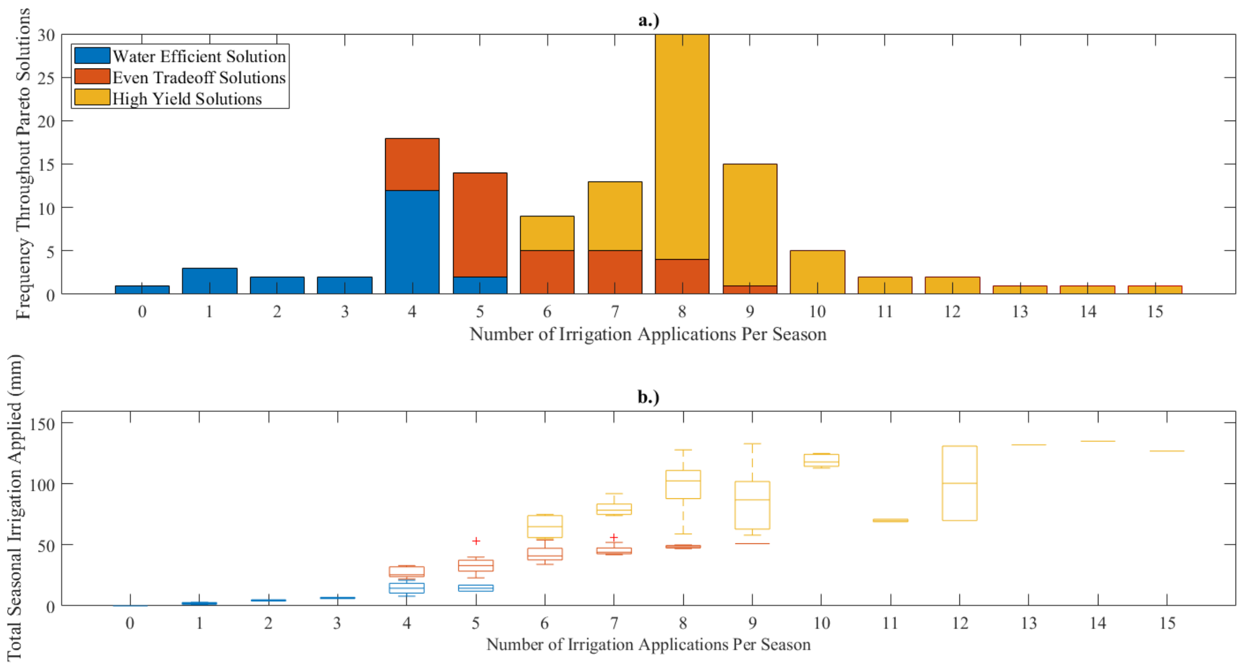

Figure 10. Application frequency analysis a) Histogram of irrigation application count among *Strategy 4* solutions. Solutions are color-coded according to their respective yield cluster (see section 2.4). b) Box plots of the total irrigation applied in *Strategy 4* solutions, divided by the number of applications applied along the x-axis. The data is further divided and color coded into the clusters as described in section 2.4.

The results of the aforementioned frequency analysis on *Strategy 4* appear in Figure 10. Figure 10.a presents *Strategy 4* application counts in the histogram format. The histogram at first appears to be bi-modal overall, but each of the three clusters actually has its own unique mode. The high yield cluster has a mode of 8 applications, the even trade-off solution cluster has a mode of 5 applications, and the water efficient solutions have a mode of 4 applications. Again, knowing the distribution of each cluster can aid farmers in developing best irrigation management practices by choosing the right number of irrigation applications to schedule in a season. Figure 10.b summarizes the statistics of the total amount of irrigation per season for each of the histogram buckets from Figure 10.a as boxplots. As Figure 10.b demonstrates, the variance of the water efficient solutions increases with the number of applications, and the variance of the even trade off cluster increases and then decreases as application count increases. Finally, the variance of the high yield solutions also initially increases with the number of applications, but then sharply decreases after 13 applications.

The histogram in Figure 11 describes the total irrigation (upper figure) and nitrogen (lower figure) application count across the three clusters from *Strategy 5*. The irrigation histogram for *Strategy 5* is also bi-modal, with peaks at 4 applications and 6 applications. Individually, the high cluster has a mode at 7 applications, the even trade-off solutions have a mode of 4 applications, and the water efficient solution has a mode of 3 applications. This central tendency towards solutions between 4 and 6 suggests that farmers restricted to 6 applications of irrigation in a season can still achieve high yielding solutions. This is another example of how EMO algorithms can help with developing management practices. The distribution of nitrogen applications is a bit more complicated. The histogram appears truncated at 6 days and cuts out of at the mode of the distribution. This suggests that Pareto optimal solutions with more than 6 days of nitrogen application exist, and in future optimization routines, more than six nitrogen application variables may yield more optimal solutions. Finding more optimal solutions beyond 6 nitrogen application variables may develop new, better management practices for maize growth in the future.

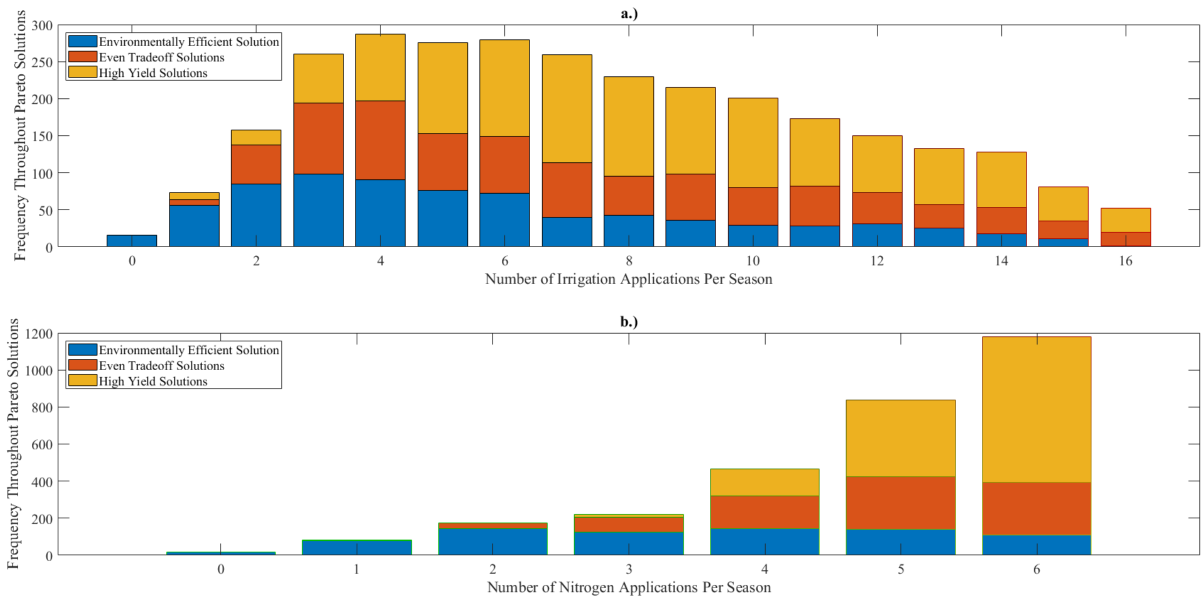

Figure 11. a.) Histogram of irrigation application count among *Strategy* 5 solutions. b.) Histogram of nitrogen application count among *Strategy* 5 solutions. Solutions are color coded according to their respective yield cluster (see section 2.4).

5. CONCLUSIONS

The EMO algorithm U-NSGA-III is thus proven to be a powerful and useful tool in developing agricultural best management practices. S*trategy 1* can provide a wide Pareto front of possible management practices, which not only work contain optima by yield, but also optima by water usage. Therefore, decision-makers can find the exact optimum for their particular agricultural use case. *Strategy 2* further refined the results from *Strategy 1* by optimizing both irrigation and nitrogen application at the same time, which better equips decision makers with management practices for their particular access to irrigation and nitrogen supplies. *Strategy 3* also added an objective that directly affects the environment, leaching.

Meanwhile, bi-level optimization runs (*Strategies 4* through *5*) built upon the first three strategies by searching beyond the fixed set of days and found more convergent and diverse population. The bi-level optimization technique also provided more solutions, which enrich the options of decision makers. The bi-level method also broke out of local optima that thwarted *Strategies 1* through *3.* Finally, examining the Pareto results of the bi-level results produces useful knowledge on the nature of application counts and their effect on other objectives.

One major limiting factor to note on this paper is the weather. The approximately 50% reduction in irrigation from *Strategy 1* appears to be an extraordinary reduction in irrigation (Table 2), but these results present the optimal obtainable output possible in that given season. In order to generate viable solutions for farmers on the ground, a decision support tool would predict the weather of an incoming season, and the tool would then run a robustness analysis on the Pareto solutions generated. In MO, robustness describes as how sensitive a solution is to slight changes in variables such as weather (Deb and Gupta, 2005). In future research papers, prioritizing robust solutions that perform well over different weather scenarios would yield more realistic Pareto optimum sets for a given season.

Finally, optimizing agricultural practices entails optimizing conflicting objectives, such as total irrigation and total yield, against each other. EMO is a powerful tool that searches for sets, or Pareto fronts, of optima within conflicting objective space. Decision makers (e.g., farmers, policymakers) can then choose the particular optimum that offers the best trade-off in their use case. The EMO algorithm U-NSGA-III, using the crop-modeling suite DSSAT as the optimization objective function, found innovative and novel solutions to irrigation and nitrogen fertilizer application planning. In addition, a novel bi-level optimization scheme (U-NSGA-III and Monte Carlo optimization) mitigates the difficulties of the unmanageably large variable space.

6. CURRENT FINDINGS AND FUTURE RESEARCH

In the coming century, societies must dramatically increase food production without furthering the irreversible damage to the environment. One means of balancing these two objectives is sustainable intensification. Sustainable intensification entails increasing the yield of existing agricultural lands while minimizing negative impacts on the environment. With these goals in mind, we sought to optimize irrigation and fertilizer scheduling on the farm level with respects to crop yield and environmental impact. Since no solution exists that maximizes yield and minimizes environmental impact, multi-objective optimization techniques were used to obtain a set of optimal solutions that collectively represent the tradeoffs between the conflicting objectives. Decision makers can then rank and select their optimal trade-off from the global set of optimal solutions. The following are the main conclusions from this research:

- Multi-objective optimization can help with sustainable agricultural intensification process by identifying solutions to irrigation scheduling and nitrogen application that are significantly more efficient than current best practices.
- Multi-objective optimization gives the power of prioritizing conflicting objectives into the hands of human decision makers by providing the Pareto front of optimal solutions
- Examining the common traits among a group of Pareto optimal solutions allows producers to improve their existing management practices
- Our bi-level optimization framework further improves the efficiency of solutions, as well as the number of Pareto optimal solutions by introducing the irrigation and nitrogen application dates in addition to the application amounts

7. FUTURE APPLICATIONS

This study integrated U-NSGA-III based multi-objective optimization platform with DSSAT crop model. The objective function of U-NSGA-III enables the platform to optimize against a myriad of soil, crop, and climate types. With our platform, we were able to produce solutions that maximize yield while reducing water usage by 48.48%, nitrogen usage by 26.4%, and nitrogen leaching by 51.48%. Despite these promising results, more work needs to be done to improve the over processes in agricultural intensification. The following are areas that can be expanded from this work in future research:

- The current research only used the climate information from one growing season. However, by quantifying the uncertainty of weather, future optimization will be able to find solutions that are more robust and applicable in the field
- Running optimization under different conditions (e.g., management practices, crop types, soils) can identifying deeper universal characteristics among optimal agricultural solutions
- Incorporating economic objectives (e.g., net profit) and economic uncertainty (e.g., market costs, materials costs) can help producers to identify more feasible solutions
- A public online decision support tool that provides producers with highly accurate day-to-day management recommendations against multiple agricultural objectives that can be a game changer by closing the yield gaps

APPENDIX

Appendix

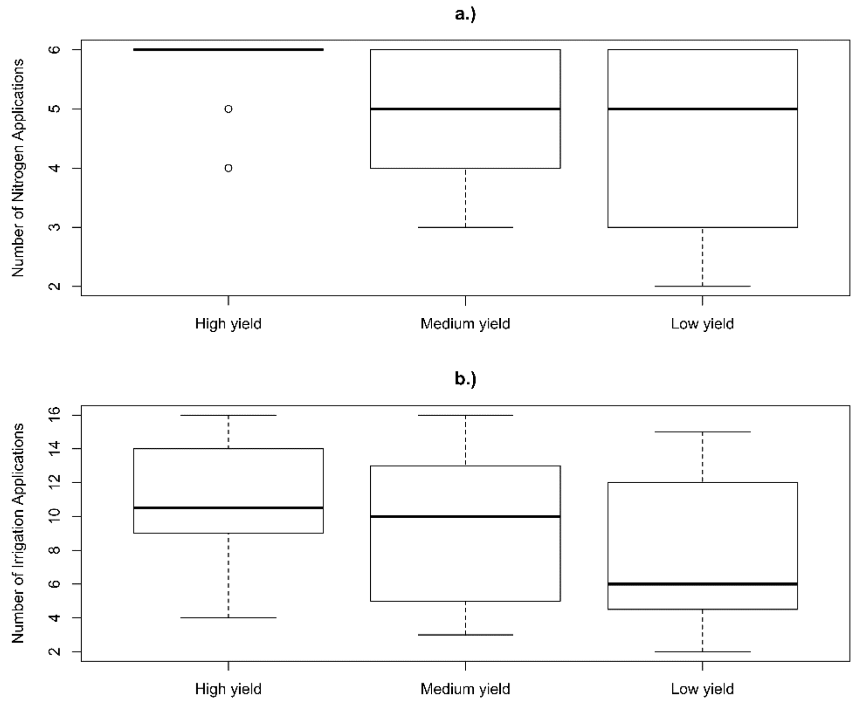

Figure A.1. Cluster analysis for a) Nitrogen application counts and b) irrigation application counts

REFERENCES

REFERENCES

30-yr Normal Maximum Temperature: Annual [WWW Document], 2015. . Prism Clim. Group, Oregon State Univ.

30-yr Normal Mean Temperature: Annual, 2015. . Prism Clim. Group, Oregon State Univ.

30-yr Normal Minimum Temperature: Annual, 2015. . Prism Clim. Group, Oregon State Univ.

30-yr Normal Precipitation: Annual [WWW Document], 2015. . Prism Clim. Group, Oregon State Univ. URL http://prism.oregonstate.edu

Akbari, M., Gheysari, M., Mostafazadeh-Fard, B., Shayannejad, M., 2018. Surface irrigation simulation-optimization model based on meta-heuristic algorithms. Agric. Water Manag. 201, 46–57.

Anwar, A.A., Haq, Z.U., 2013. Genetic algorithms for the sequential irrigation scheduling problem. Irrig. Sci. 31, 815–829. https://doi.org/10.1007/s00271-012-0364-y

Arrow, K., 1951. Social Choice and Individual Values John Wiley and Sons New York Google Scholar.

Arthur, D., Vassilvitskii, S., 2007. k-means++: The advantages of careful seeding, in: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 1027–1035.

Ashofteh, P.-S., Haddad, O.B., Loáiciga, H.A., 2015. Evaluation of Climatic-Change Impacts on Multiobjective Reservoir Operation with Multiobjective Genetic Programming. J. Water Resour. Plan. Manag. 141, 04015030. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000540

Ashrafi, S.M., Dariane, A.B., 2013. Performance evaluation of an improved harmony search algorithm for numerical optimization: Melody Search (MS). Eng. Appl. Artif. Intell. 26, 1301–1321. https://doi.org/10.1016/j.engappai.2012.08.005

Baker, D.N., Lambert, J.R., McKinion, J.M., 1983. GOSSYM: A simulator of cotton crop growth and yield. South Carolina. Agric. Exp. Station. Tech. Bull.

Beni, G., Wang, J., 1993. Swarm intelligence in cellular robotic systems, in: Robots and Biological Systems: Towards a New Bionics? Springer, pp. 703–712.

Benson, H.P., 1978. Existence of efficient solutions for vector maximization problems. J. Optim. Theory Appl. 26, 569–580. https://doi.org/10.1007/BF00933152

Boote, K.J., Jones, J.W., Hoogenboom, G., 1998. Simulation of crop growth: CROPGRO model.

Boote, K.J., Jones, J.W., Singh, P., 1992. Modeling growth and yield of groundnut.

Chung, U., Gbegbelegbe, S., Shiferaw, B., Robertson, R., Yun, J.I., Tesfaye, K., Hoogenboom, G., Sonder, K., 2014. Modeling the effect of a heat wave on maize production in the USA and its implications on food security in the developing world. Weather Clim. Extrem. 5, 67–77. https://doi.org/10.1016/j.wace.2014.07.002

Consoli, S., Matarazzo, B., Pappalardo, N., 2008. Operating Rules of an Irrigation Purposes Reservoir Using Multi-Objective Optimization. Water Resour. Manag. 22, 551–564. https://doi.org/10.1007/s11269-007-9177-9

Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C., 2009. Introduction to Algorithms. MIT Press. https://doi.org/10.1108/S0749-742320160000019022

Cox, W., While, L., 2016. Improving the IWFG algorithm for calculating incremental hypervolume, in: 2016 IEEE Congress on Evolutionary Computation (CEC). pp. 3969–3976. https://doi.org/10.1109/CEC.2016.7744293

de Borda, J., 1781. Mathematical derivation of an election system. Isis 44, 42–51.

Deb, K., 2009. Multi-objective Optimization Using Evolutionary Algorithms. John Wiley and Sons, Chichester, UK.

Deb, K., Agrawal, R.B., 1995. Simulated Binary Crossover for Continuous Search Space. Complex Syst. 9, 115–148. https://doi.org/10.1.1.26.8485Cached

Deb, K., Gupta, H., 2005. Searching for Robust Pareto-Optimal Solutions in Multi-objective Optimization, in: Coello Coello, C.A., Hernández Aguirre, A., Zitzler, E. (Eds.), Evolutionary Multi-Criterion Optimization. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 150–164.

Deb, K., Jain, H., 2014. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints. IEEE Trans. Evol. Comput. 18, 577–601. https://doi.org/10.1109/TEVC.2013.2281535

Deb, K., Pratap, A., Agarwal, S., Meyarivan, T., 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans. Evol. Comput. 6, 182–197. https://doi.org/10.1109/4235.996017

Eberhart, R., Kennedy, J., 1995. A new optimizer using particle swarm theory. MHS'95. Proc. Sixth Int. Symp. Micro Mach. Hum. Sci. 39–43. https://doi.org/10.1109/MHS.1995.494215

Emmerich, M., Beume, N., Naujoks, B., 2005. An EMO Algorithm Using the Hypervolume Measure as Selection Criterion, in: Coello Coello, C.A., Hernández Aguirre, A., Zitzler, E. (Eds.), Evolutionary Multi-Criterion Optimization. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 62–76.

Fernández García, I., Moreno, M.A., Rodríguez Díaz, J.A., 2014. Optimum pumping station management for irrigation networks sectoring: Case of Bembezar MI (Spain). Agric. Water Manag. 144, 150–158. https://doi.org/10.1016/j.agwat.2014.06.006

García-Vila, M., Fereres, E., Mateos, L., Orgaz, F., Steduto, P., 2009. Deficit Irrigation Optimization of Cotton with AquaCrop. Agron. J. 101, 477. https://doi.org/10.2134/agronj2008.0179s

Geem, Z.W., Kim, J.H., Loganathan, G. V., 2007. A New Heuristic Optimization Algorithm: Harmony Search. https://doi.org/10.1177/003754970107600201

Gen, M., Cheng, R., 2000. Genetic algorithms and engineering optimization. John Wiley & Sons.

Geoffrion, A.M., 1968. Proper efficiency and the theory of vector maximization. J. Math. Anal. Appl. 22, 618–630. https://doi.org/10.1016/0022-247X(68)90201-1

Goldberg, D.E., 1989. Genetic algorithms in search, optimization, and machine learning. Addison-Wesley, Boston.

González Perea, R., Camacho Poyato, E., Montesinos, P., Rodríguez Díaz, J.A., 2016. Optimization of Irrigation Scheduling Using Soil Water Balance and Genetic Algorithms. Water Resour. Manag. 30, 2815–2830. https://doi.org/10.1007/s11269-016-1325-7

Groot, J.C.J., Oomen, G.J.M., Rossing, W.A.H., 2012. Multi-objective optimization and design of farming systems. Agric. Syst. 110, 63–77. https://doi.org/10.1016/j.agsy.2012.03.012

Groot, J.C.J., Rossing, W.A.H., Jellema, A., Stobbelaar, D.J., Renting, H., Van Ittersum, M.K., 2007. Exploring multi-scale trade-offs between nature conservation, agricultural profits and landscape quality-A methodology to support discussions on land-use perspectives. Agric. Ecosyst. Environ. 120, 58–69. https://doi.org/10.1016/j.agee.2006.03.037

Gurav, J.B., Regulwar, D.G., 2012. Multi Objective Sustainable Irrigation Planning with Decision Parameters and Decision Variables Fuzzy in Nature. Water Resour. Manag. 26, 3005–3021. https://doi.org/10.1007/s11269-012-0062-9

Hadka, D., Reed, P., 2013. Borg: An Auto-Adaptive Many-Objective Evolutionary Computing Framework 21, 231–259.

Haimes, Y. V, 1971. On a bicriterion formulation of the problems of integrated system identification and system optimization. IEEE Trans. Syst. Man. Cybern. 1, 296–297.

Hengsdijk, H., Langeveld, J.W.A., 2009. Yield trends and yield gap analysis of major crops in the world.

Holland, J.H., 1975. Adoption in Natural and Artificial Systems. University of Michigan. Ann Arbor.

Hoogenboom, G., Jones, J.W., Traore, P.C.S., Boote, K.J., 2012. Experiments and data for model evaluation and application, in: Improving Soil Fertility Recommendations in Africa Using the Decision Support System for Agrotechnology Transfer (DSSAT). Springer, pp. 9–18.

Hoogenboom, G., Porter, C.H., Shelia, V., Boote, K.J., Singh, U., White, J.W., Hunt, L.A., Ogoshi, R., Lizaso, J.I., Koo, J., 2017. Decision Support System for Agrotechnology Transfer (DSSAT) Version 4.7.

Jones, C.A., Kiniry, J.R., Dyke, P.T., 1986. CERES-Maize: A simulation model of maize growth and development. Texas A& M University Press.

Jones, J., Hoogenboom, G., Porter, C., Boote, K., Batchelor, W., Hunt, L., Wilkens, P., Singh, U., Gijsman, A., Ritchie, J., 2003. The DSSAT cropping system model. Eur. J. Agron. 18, 235–265. https://doi.org/10.1016/S1161-0301(02)00107-7

Karami, F., Dariane, A.B., 2018. Many-Objective Multi-Scenario Algorithm for Optimal Reservoir Operation Under Future Uncertainties. Water Resour. Manag. 32, 3887–3902. https://doi.org/10.1007/s11269-018-2025-2

Kijne, J.W., Barker, R., Molden, D.J., 2003. Water productivity in agriculture: limits and opportunities for improvement. Cabi.

Kirkby, M., 1975. Hydrograph modelling strategies. IN Peel, R., Chisholm, M. & Haggett, P.(Eds.) Processes in Physical and Human Geography.

Kollat, J.B., Reed, P.M., 2006. Comparing state-of-the-art evolutionary multi-objective algorithms for long-term groundwater monitoring design. Adv. Water Resour. 29, 792–807. https://doi.org/10.1016/J.ADVWATRES.2005.07.010

Koza, J.R., 1992. Genetic programming: On the programming of computers by means of natural selection. MA.

Kropp, I., 2018a. DSSAT4j.

Kropp, I., 2018b. CropOpt.

Lampinen, J., Zelinka, I., others, 2000. On stagnation of the differential evolution algorithm, in: Proceedings of MENDEL. pp. 76–83.

Lascody, R., 2002. The onset of the wet and dry seasons in east central Florida-A subtropical wet-dry climate? Natl. Ocean. Atmos. Adminstration/National Weather Serv. Forecast Off. Melbourne, FL.

Laumanns, M., Thiele, L., Deb, K., Zitzler, E., 2002. Combining Convergence and Diversity in Evolutionary Multiobjective Optimization. Evol. Comput. 10, 263–282. https://doi.org/10.1162/106365602760234108

Mello Jr, A. V, Schardong, A., Pedrotti, A., 2013. Multi-Objective Analysis Applied to an Irrigated Agricultural System on Oxisols. Trans. ASABE 56, 71–82.

Miettinen, K., 2012. Nonlinear multiobjective optimization. Springer Science & Business Media.

Nixon, J.B., Dandy, G.C., Simpson, A.R., 2001. A genetic algorithm for optimizing off-farm irrigation scheduling. J. Hydroinformatics 3, 11–22.

Quinn, J.D., Reed, P.M., Giuliani, M., Castelletti, A., Oyler, J.W., Nicholas, R.E., 2018. Exploring How Changing Monsoonal Dynamics and Human Pressures Challenge Multireservoir Management for Flood Protection, Hydropower Production, and Agricultural Water Supply. Water Resour. Res. 54, 4638–4662.

Ritchie, J.T., 1985. Description and performance of CERES wheat: A user-oriented wheat yield model. ARS wheat yield Proj. 159–175.

Rudolph, G., 1994. Convergence Analysis of Canonical Genetic Algorithms. IEEE Trans. Neural Networks 5, 96–101. https://doi.org/10.1109/72.265964

Sampath, R.K., 1992. Issues in irrigation pricing in developing countries. World Dev. 20, 967–977. https://doi.org/10.1016/0305-750X(92)90124-E

Sarker, R., Ray, T., 2009. An improved evolutionary algorithm for solving multi-objective crop planning models. Comput. Electron. Agric. 68, 191–199. https://doi.org/10.1016/j.compag.2009.06.002

Sarker, R., Ray, T., 2005. Multiobjective Evolutionary Algorithms for solving Constrained Optimization Problems.

Sastry, K., Goldberg, D.E., Kendall, G., 2014. Genetic Algorithms BT - Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques, in: Burke, E.K., Kendall, G. (Eds.), . Springer US, Boston, MA, pp. 93–117. https://doi.org/10.1007/978-1-4614-6940-7_4

Seada, H., 2017. evolib.

Seada, H., Deb, K., 2015. U-NSGA-III: A Unified Evolutionary Optimization Procedure for Single, Multiple, and Many Objectives: Proof-of-Principle Results. Springer, Cham, pp. 34–49. https://doi.org/10.1007/978-3-319-15892-1_3

Seada, H., Kropp, I., 2018. evolib (multi-threaded).

Smith, M., 1992. CROPWAT: A computer program for irrigation planning and management. Food & Agriculture Org.

Steduto, P., Hsiao, T.C., Raes, D., Fereres, E., 2009. Aquacrop-the FAO crop model to simulate yield response to water: I. concepts and underlying principles. Agron. J. 101, 426–437. https://doi.org/10.2134/agronj2008.0139s

Storn, R., Price, K., 1997. Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. J. Glob. Optim. 11, 341–359. https://doi.org/10.1023/a:1008202821328

Tamaki, H., Kita, H., S, K., 1996. Multi-objective optimization by genetic algorithms: a review. Diversity. https://doi.org/10.1080/02680939.2013.782512

Tan, Q., Zhang, S., Li, R., 2017. Optimal use of agricultural water and land resources through reconfiguring crop planting structure under socioeconomic and ecological objectives. Water (Switzerland) 9, 488. https://doi.org/10.3390/w9070488

Wang, Y., Wu, P., Zhao, X., Jin, J., 2012. Water-Saving Crop Planning Using Multiple Objective Chaos Particle Swarm Optimization for Sustainable Agricultural and Soil Resources Development. CLEAN – Soil, Air, Water 40, 1376–1384. https://doi.org/10.1002/clen.201100310

Waraich, E.A., Ahmad, R., Ashraf, M.Y., Saifullah, Ahmad, M., 2011. Improving agricultural water use efficiency by nutrient management in crop plants. Acta Agric. Scand. Sect. B-Soil Plant Sci. 61, 291–304.

Whittaker, G., Färe, R., Grosskopf, S., Barnhart, B., Bostian, M., Mueller-Warrant, G., Griffith, S., 2017. Spatial targeting of agri-environmental policy using bilevel evolutionary optimization. Omega (United Kingdom) 66, 15–27. https://doi.org/10.1016/j.omega.2016.01.007

Wilkerson, G.G., Jones, J.W., Boote, K.J., Ingram, K.T., Mishoe, J.W., 1983. Modeling soybean growth for crop management. Trans. ASAE 26, 63–73.

Xue, F., Sanderson, A.C., Graves, R.J., 2003. Pareto-based multi-objective differential evolution. 2003 Congr. Evol. Comput. CEC 2003 - Proc. 2, 862–869. https://doi.org/10.1109/CEC.2003.1299757

Zhang, C., Li, Y., Chu, J., Fu, G., Tang, R., Qi, W., 2017. Use of Many-Objective Visual Analytics to Analyze Water Supply Objective Trade-Offs with Water Transfer. J. Water Resour. Plan. Manag. 143, 5017006.

Zhang, Q., Li, H., 2007. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. IEEE Trans. Evol. Comput. 11, 712–731.