Augmented Evolutionary Intelligence: Combining Human and Evolutionary Design for Water Distribution Network Optimisation

Matthew B. Johns
University of Exeter
Exeter, EX4 4QF
United Kingdom
M.B.Johns@exeter.ac.uk

Herman A. Mahmoud
University of Exeter
Exeter, EX4 4QF
United Kingdom
H.Mahmoud@exeter.ac.uk

David J. Walker
University of Plymouth
Plymouth, PL4 8AA
United Kingdom
david.walker@plymouth.ac.uk

Nicholas D. F. Ross
University of Exeter
Exeter, EX4 4QF
United Kingdom
nr339@exeter.ac.uk

Edward C. Keedwell
University of Exeter
Exeter, EX4 4QF
United Kingdom
E.C.Keedwell@exeter.ac.uk

Dragan A. Savic
University of Exeter/KWR
Exeter/Nieuwegein
United Kingdom/Netherlands
D.Savic@exeter.ac.uk

ABSTRACT

Evolutionary Algorithms (EAs) have been employed for the optimisation of both theoretical and real-world problems for decades. These methods although capable of producing near-optimal solutions, often fail to meet real-world application requirements due to considerations which are hard to define in an objective function. One solution is to employ an Interactive Evolutionary Algorithm (IEA), involving an expert human practitioner in the optimisation process to help guide the algorithm to a solution more suited to real-world implementation. This approach requires the practitioner to make thousands of decisions during an optimisation, potentially leading to user fatigue and diminishing the algorithm's search ability. This work proposes a method for capturing engineering expertise through machine learning techniques and integrating the resultant heuristic into an EA through its mutation operator. The human-derived heuristic based mutation is assessed on a range of water distribution network design problems from the literature and shown to often outperform traditional EA approaches. These developments open up the potential for more effective interaction between human expert and evolutionary techniques and with potential application to a much larger and diverse set of problems beyond the field of water systems engineering.

CCS CONCEPTS

• Theory of computation → Design and analysis of algorithm

KEYWORDS

Evolutionary Algorithm, Machine Learning, Human-computer Interaction, Knowledge Guided Search, Water Distribution Network Design, Real-world Application

1 INTRODUCTION

Evolutionary Algorithms (EAs) have been used for the optimization of a wide range of both theoretical and real-world problems spanning many fields, including water systems engineering. It is often difficult to generate solutions using an EA that are suitable for real-world application in this field without substantial intervention on the part of an expert. This is due in part to the difficulty of defining every consideration an engineer needs to take into account when designing a complex system such as a water distribution network (WDN) and expressing this mathematically as an objective function or constraint.

This paper proposes a method for the capture and integration of engineering expertise into an EA through the use of interactive visualization and machine learning techniques with the aim to create more optimal and engineering feasible solutions from an EA.

1.1 Knowledge Guided Search

The EA has proven to be a versatile process for solving a large variety of optimization problems spanning many fields and disciplines. The strength of the approach comes from the ability it has to traverse large search spaces, avoiding local optima and therefore can be viewed as a truly global search technique. The performance and versatility of the EA can be attributed partly to the independence it has over the problem being undertaken. Although seen as an asset, this problem independence can have a detrimental effect on performance in the case where the algorithm has not been tuned to a great enough extent to solve the problem at hand.

For the problem of WDN design the EA often relies on operators such as crossover and mutation to alter the configuration of the network. These operators however are blind to the direct effect any changes made to elements of the network have on the overall performance of the resultant solution. For example, from the perspective of the EA, a change in the diameter of a pipe has no

bearing on the hydraulic behavior of connected elements until the resultant design is evaluated, although an engineer would know that the pressure at adjacent junctions would be affected. The performance of a newly created network is only known following solution decoding and hydraulic simulation and although this abstraction is partly why EAs can be applied to many different water system design problems, there is clear scope for the integration of problem specific knowledge, something that has started to be explored in the literature [1]–[3].

There have been several approaches in the literature which use knowledge of the problem or search space to aid the search of an EA. One such method is Guided Local Search (GLS) [4], a metaheuristic technique bearing similarities to tabu search and simulated annealing. GLS has displayed good performance on a number of combinatorial optimization problems [5]–[8] as it helps prevent the search from becoming stuck in a local minima. GLS functions by penalizing certain solution features that it deems would not occur in near optimal solutions through the use of weighted penalties. Another approach is Guided Mutation (GM) [9], an offspring generating operator for EAs which is considered a combination between an EA's standard mutation operator and the offspring generating method of an Estimation of Distribution Algorithm [10]. GM works on the notion that good solutions have a similar structure and new offspring should be generated close to the good solutions already found during the search and are fully automated techniques.

1.2 Interactive Evolution

Interactive Evolution (IE) [11] aims to incorporate knowledge through human interaction with an EA which requires input from a user during the search process. User interaction is commonly used to assess a solution's fitness; however, the user can also be involved during the variation and selection phases of the algorithm. A common issue when applying an EA to a problem, especially in a real-world setting, is there are often non-explicit conditions that are very difficult to define. Various design approaches require the human to make subjective decisions based on human intuition, such as the ability to judge a design's aesthetic qualities in the case of art or furniture design [12]. The fitness criterion in cases such as these cannot be explicitly formulated and often require case-by-case comparison to effectively assess a solution. The interaction of a human user can also be employed to more effectively guide an algorithm's search of the solution space with the view to speed up convergence and prevent local optima trappings.

Water resources design and management problems are complex to solve; not only from a mathematical perspective, but also from political, sociological, and other subjective viewpoints. The majority of research in the field of water resources concentrates on the improvement of simulation models and their incorporation with optimization techniques such as evolutionary algorithms. The problem lies in real life cases where the optimization technique employed returns a mathematically optimal solution, however the solution may become infeasible when considering subjective preferences [13]. Recently, researchers in the field of water resources have developed methods for the calibration of models

through the use of interactive evolution which enables the incorporation of unmodeled objectives in the search procedure [14]. The field of interactive evolution is a rapidly growing area of research; with the aim to utilize the subjective responses from human users to guide the search of evolutionary algorithms [11]. Singh et al. [14] used an elitist non-dominated sorting genetic algorithm (NSGA-II) [15] and human responses to find optimal solutions for groundwater problems which were both mathematically optimal and feasible. This was achieved through the consideration of human responses as one of the multiple criteria for the computation of the solution fitness. Although the interaction element of the process was simple (solution ranking) the results of the study were successful in generating superior solutions than non-interactive optimization runs of NSGA-II.

1.3 Machine Learning

Any approach that seeks to embed human expertise into an EA must have a mechanism to learn from the user. In previous research this has been achieved through the expression of 'rules of thumb' which are embedded into the EA through heuristics. However, this is difficult to achieve in reality as most decisions made by an expert user will be based on intuition and 'feel' rather than explicit rules. Therefore, in this paper we introduce the use of machine learning as a mechanism to learn user behavior from interaction and to embed knowledge within the EA. Machine learning is a data analytics technique that teaches computers to achieve what comes naturally to humans and learn from experience [16]. Machine learning algorithms use computational methods to learn knowledge directly from data without depending on equations as model. These algorithms, in particular, neural networks and decisions trees have been used in water distribution network applications such as leakage detection, calibration models, demand forecasting models and pipe deterioration modelling. A decision tree algorithm, used in this work, describes a class of methods to cope with model classification and regression problems in machine learning. The main benefits of decision trees are their simplicity and computational efficiency, both in terms of creating the tree as well as applying it to decision-making [17]. Also, decision trees are easy to understand, able to handle large data, and the resulting trees can directly be visualized and interpreted.

1.4 Multi-objective Water Distribution Network Design

Different criteria such as cost, reliability and water quality are used to find the optimal solutions in real-world WDN design problems. Thus, many existing studies focus on multi-objective (in particular two-objective) optimal design of WDN. The first objective normally involves minimizing the total cost of the network, whilst the second objective involves maximization of network benefits [18].

In this work, the total cost (i.e., first objective criteria), denoted by *CN (Eq. 1),* includes the initial capital expenditure for pipes. The network benefits (i.e., second objective criteria) is measured through network resilience index (*RI*) (Eq. 2). *RI* has been shown to be a promising measure in comparison to others available

measures in the literature [19]. Further, *RI* considered both excess pressure head at each demand node and the uniformity of pipes connected to that demand node. Thus, the operation of a network with more pressure than a consumer requires (i.e., surplus pressure) can provide additional resilience in the event of failure such as bursts or fire flows. The above dual-objective problem can be expressed mathematically as follows:

$$Min \qquad CN = \sum_{k=1}^{np} U_k(D_k) \times L_k \qquad (1)$$

$$Max \qquad RI = \frac{\sum_{n=1}^{nn} C_n Q_n \left(P_n^{avl} - P_n^{req}\right)}{\left(\sum_{r=1}^{nr} Q_r P_r + \sum_{p=1}^{npu} \frac{P_p}{\gamma}\right) - \sum_{n=1}^{nn} Q_n^{avl} P_n^{req}} \qquad (2)$$

$$C_n = \frac{\sum_{k=1}^{npn} D_{kn}}{npn \times \max\{D_{kn}\}} \qquad (3)$$

Where $U_k(D_k)$= unit cost of a pipe of a given diameter; L_k=length of pip*e k; np*= number of pipes in a given network*; nn*= number of demand nodes in a given network; Q_n= demand at node *n*; P_n^{avl} = available pressure head at node *n*; P_n^{req} = required pressure head at node *n*; Q_r= supply at reservoir *r*; P_r= elevation head at reservoir *r*; *nr*= number of reservoirs in given network; *npu*= number of pumps in a given network; P_p=power of pump *p; γ*= specific weight of wat*er;* C_n= uniformity at node *n*; *npn*= the number of pipes connected to node *n;* and D_{kn} = the diameter of pipe *k* connected to node *n*.

A generated solution to the above problem is represented by a vector of integers, in which each element is the value of a pipe diameter in that solution. The value of the vector ranges from one up to the number of commercially available diameters. The quality of the generated solution is evaluated under two objective functions given in Eqs. 1 and 2.

The above design problem is subject to the hydraulic constraints. These involve satisfying continuity at each demand node, conserving energy in loops and ensuring that available pressure head at each node is always equal to or above the required pressure head. The above constraints require solving conservation of mass and energy equations to determine the nodal pressure heads, flows in pipes for a given network, and are automatically satisfied by using the well-known EPANET2.0 [20] hydraulic solver.

Three benchmark WDNs from the literature were selected to assess the capabilities of Human Derived Heuristic (HDH) based methods presented in this paper. The networks range in size, complexity, and network features providing different levels of challenge for both engineer and algorithm. The first WDN is Hanoi [21], a representation of s single water source (i.e. reservoir) network consisting of three loops, based upon the trunk main layout for the city of Hanoi, Vietnam. It consists of 34 decision pipes and 6 available pipe diameters. The second test network is Blacksburg [22], a representation of a single source network consisting of multiple loops and branches. It consists of 35 pipes 23 of which are decision variables and 14 available pipe diameters. The final WDN

on test is Modena [23], a representation of the water supply system of the city of Modena, Italy. The network consists of four sources and 317 decision pipes with 13 available pipe diameters to choose from.

2 EXPERIMENTAL SETUP

The experimentation presented in this paper is comprised of three separate parts; interaction capture, human-derived heuristic learning, and the integration of human-derived heuristics into EAs. As previously stated, the aim of this work is to develop a method for integrating expert engineering knowledge into an EA with the ultimate view of creating an algorithm which automatically learns from its interactions with a human expert.

The following approach was demonstrated on three WDN design problems of varying size and complexity from the literature; Hanoi, Blacksburg and Modena.

2.1 Engineering Interaction Capture

A software framework for WDN optimization (HOWS framework[24], [25]) was developed for the interactive optimization of WDN design and operation problems. The framework employs a server-client architecture, where the server manages the configuration and automatic optimization of the problem and the client is tasked with visualization and user interaction capture operations. The client presents the user with a three-dimensional representation of a WDN. Various visualization techniques are used to convey topographical, hydraulic, and optimization information to aid the user in decision making. Figure 1 shows the Blacksburg network in the interactive visualization client. The most prevalent components in a WDN are junctions and pipes, these are represented as spheres and cylinders respectively. In this configuration of the client, the diameter of a cylinder is proportional to the diameter of its respective pipe, enabling the user to quickly identify diameters without necessarily interactively inspecting the pipe in question. The network topology is primarily defined by the position of the junctions in 3D space providing the engineer with an idea of distance and elevation change throughout the network. In this configuration, the hydraulic head deficit values are show at each junction using a linear color scale where green indicates head constraint satisfaction and yellow to red (red being maximum head deficit) specifies varying degrees of head constraint violation, allowing the user to quickly identify areas of the network violating problem constraints.

Figure 1: HOWS Framework Interactive Visualization Client

Additional hydraulic information is displayed to the user with 'pipe fins', these run the length of each pipe and can be used to show a variety of variables simultaneously. For this set of experiments, the height at each end represents hydraulic head at the adjoining junctions, this value is reinforced using a color gradient where green is high, and red is low. The vertical lines which can be seen running the length of the fins are moved in the direction of flow with a speed relative to velocity of water in the pipe. The inclusion of the pipe fins in this experiment allows the engineer to gain greater insight into the performance of a network with the view to helping them make more informed decisions without the need to constantly alter the visualization settings, which could potentially slow the decision-making process and reduce collection volume of interaction data.

For this set of experiments the user is first presented with a randomly generated solution to the problem and instructed to manually optimize the network, reducing network infrastructure cost whilst meeting the basic pressure constraints of the problem. These values are displayed to the user in the top left of the screen. The user interacts with the network by clicking on the components of the network. Selecting a junction will display its constraint information, highlighting the amount of hydraulic head deficit or excess. Selecting a pipe will bring up a dialog displaying a list of available diameters and their associated cost. The user can then select a new diameter. After the user changes the diameter of a pipe, the change is sent to the server which logs the change and runs a hydraulic simulation on the new network configuration and computes the new objective and constraint values. This information is then sent back to the client which updates the representation of the network to reflect the updated information.

2.2 Learning Human-derived Heuristics

This work builds upon previous work [26] that used decision trees to model human knowledge when optimizing a WDN design problem. The results of that study indicated that there was scope for the capture of expert water systems knowledge and its integration into an EA. The approach however, relied heavily on training the models using the entire network state, resulting in a very large number of model features (the diameters of all pipes and pressures at every junction). The approach presented here was developed

with a view of reducing the number of model features required as a step towards developing a more generalizable method for water system knowledge-based model generation.

The task of the model presented here is to predict the diameter of a randomly selected pipe a human would choose given the network's current state. It was decided that a decision tree-based learning approach would again be employed in this work due the ability to visualize and interpret the generated models; aiding in analysis of the effectiveness of the generated models.

Decision trees require a fixed input schema. In this approach four input features, local to the selected pipe are considered; the current diameter, water velocity, upstream head deficit, and downstream head deficit. These parameters were chosen as they are thought to be the primary considerations of the engineer when selecting a new diameter. The algorithm used to generate the decision trees is an optimized version [27] of the Classification and Regression Trees (CART) algorithm [28]. Following generation, the models are assessed using explained variance and the leave-one-out cross validation method. Given a set of K recorded interactions made by the user, the model is trained on K-1 of the observations, and then tested on the remaining unseen observation. This procedure is repeated K times, such that each observation is the test case exactly once. For the first two networks, Hanoi and Blacksburg, interactions were recorded for three users, each of which optimized each network multiple times. For the Modena network one user optimized the network twice.

2.3 Integrating Heuristics into EAs

The trained and validated decision tree model is integrating into an EA through the mutation operator. The HDH mutation operator is designed to take the place of an EAs standard mutation procedure. The HDH mutation firstly decodes the chromosome and randomly selects a pipe in the network, then the selected pipe's diameter, upstream head deficit, downstream head deficit and velocity are applied to the HDH model which predicts the new diameter for the pipe. This value is compared with the available diameters for that pipe and the closest diameter is applied to the selected pipe.

An important consideration when integrating problem-specific knowledge into an EA is computational efficiency. The most computationally demanding operations are solution evaluations and in the case of water distribution design problems this comes in the form of the hydraulic simulations. Therefore, it is important not to incur any additional objective function evaluations where possible. Due to the dependency the HDH mutation operator has on a solution's pressure and velocity information, mutation cannot be applied post crossover without the need to re-evaluate the hydraulic network of resultant solutions. Therefore, the mutation operator needs to precede the crossover operator in order to preserve the hydraulic information gained from the hydraulic simulation of the original solution.

Through initial experimentation it was found that completely replacing standard bit-flip mutation with HDH mutation resulted in the premature convergence of the algorithm, therefore it was important to implement a method to control the application strength

of the HDH, essentially mixing HDH and bit-flip mutation. The first set of optimization experiments presented in this paper show the effect different application strengths of the HDH have on the algorithm search performance.

The EAs employed in this set of experiments are NSGAII and the Strength Pareto EA 2 (SPEA2) algorithm. Both algorithms have been shown to perform well on this problem type [18], [29].

3 RESULTS AND DISCUSSION

The experimental results in this paper are presented in two sections. The first being the analysis of the decision trees generated from the user interactions with each WDN problem. The final section details the application of human derived models in two EAs and the performance is assessed.

3.1 Decision Tree Analysis

For each WDN a separate decision tree was trained using the interaction data gained from the HOWS platform detailed in the previous section.

Figure 2 shows the trimmed decision tree generated from the cumulative user interaction sessions. It can be observed that in general, lower velocity values result in smaller diameters being selected. High velocity in a pipe can indicate large head-loss which reduces pressures downstream. This is something that the designer normally wants to avoid and will look to smooth out velocities in a network. The model seems to have captured some basic rules that water system engineers often employ, for example, when a pipe has high downstream deficit and upstream excess the diameter of the pipe is increased, thus working towards eliminating bottlenecks in a pipe series.

The effectiveness of the models was then assessed using the leave-one-out cross validation method using explained variance to score the model's accuracy. Table 1 shows the results of the cross-validation experiments for each network model. The explained variance metric gives an indicator as to how well the model fits the data, where 1 is 100% accuracy.

Table 1: Cross Validation Results

	Explained Variance	Standard Deviation
Hanoi	0.6	0.11
Blacksburg	0.31	0.13
Modena	0.43	0.13

The model trained on the Hanoi problem obtains the highest accuracy with the Blacksburg model achieving the lowest. The lower accuracy of Blacksburg and Modena could indicate that the features considered by the model only partially capture the decisions making of the engineering interactions. An additional consideration is that the Blacksburg and Modena problems have a much higher density of available pipe diameters (14 and 13 respectively) when compared to the Hanoi problem (6 diameters). This could go some way to explaining the reduced performance of the two models.

3.2 Optimization of WDN Using HDH

For each problem presented in this section, the base algorithm (NSGAII and SPEA2) parameters remain constant. Population size of 100, single-point crossover and a bit mutation probability = 1/n where n in the number of bits in the chromosome.

The first set of experimental results presented explore the impact varying the HDH application strength has on NSGAII's performance on the Hanoi problem. The experiment involved standard NSGAII and four variants of NSGAII with HDH mutation (NSGAII - HDH) operating with a probability of HDH application at 0.25 intervals. Each algorithm was run 30 times with a termination criterion on 500,000 fitness evaluations. Figure 3 shows the average hypervolume [30] of NSGAII and the NSGAII-HDH variants over the search. The primary finding from these results is that if the HDH is applied exclusively, the algorithm prematurely converges at a sub-optimal solution compared to the standard configuration of NSGAII. This result is perhaps expected given that the human heuristics lack the explorative capability of a standard random mutation. The remaining NSGAII - HDH variants however all ultimately perform better than NSGA-II, with a 50% probability of HDH application exhibiting the best performance.
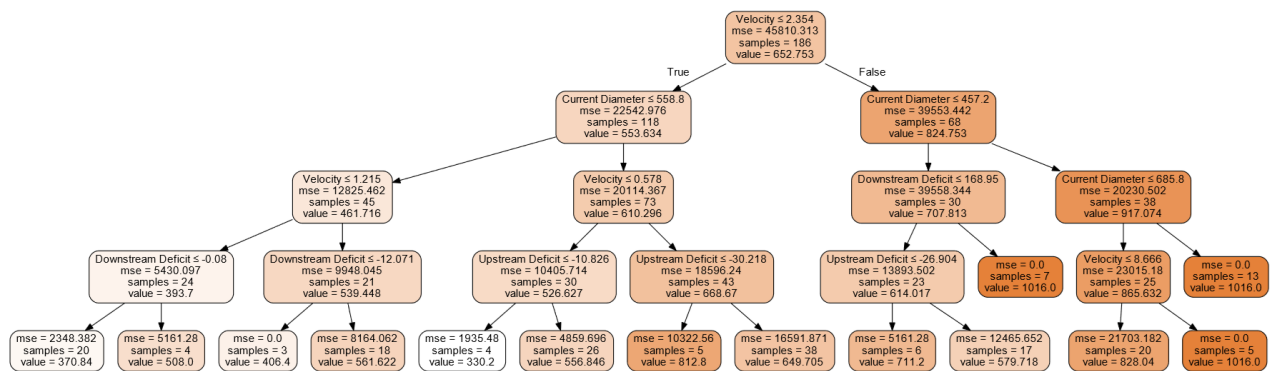

Figure 2: Decision Tree for the Hanoi Problem

The core finding of these results is that the algorithm requires some random, in this case bit-flip, mutation to fully explore the search space and prevent premature convergence.
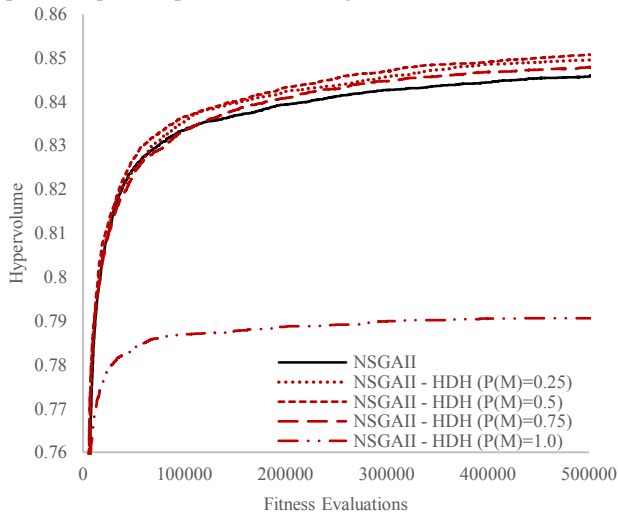

Figure 3: Mean Hypervolume for the Hanoi Problem - NSGAII - HDH Application Probability Results

Following these results, it was decided that the NSGAII variants would apply the HDH with 25% and 50% probability. Due to the decreased overall performance of SPEA2 on all test cases, HDH application strength tuning was deemed unnecessary to gauge HDH impact therefore HDH is applied with a 50% probability for all problems. Both NSGAII and SPEA2, including their HDH variants are applied to the Hanoi problem. Figure 4 shows the average hypervolume of the algorithms over 500,000 evaluations. In the case of NSGAII, the addition of the human trained model boosts performance as the search progresses, ultimately achieving a better spread of solutions. Interestingly SPEA2 only sees benefit from the HDH in the early stages of the search, with SPEA2 eventually surpassing SPEA2-HDH at approximately 120,000 evaluations.
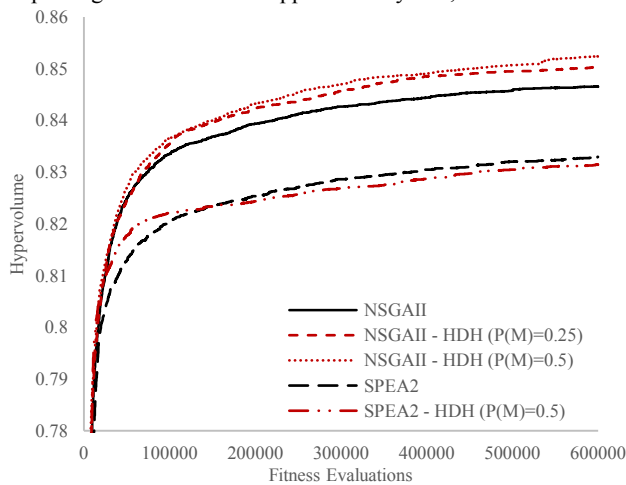

Figure 4: Mean Hypervolume for the Hanoi Problem – NSGAII, NSGAII – HDH, SPEA2 & SPEA2 – HDH

Figure 5 shows the final hypervolume results (30 runs) for each of the algorithms on test. Comparing the NSGAII results shows the human trained model has a positive impact on the solution quality of the algorithm's final populations. Not only is mean hypervolume increased but variance is drastically decreased in the case of 50% HDH mutation. The difference between the SPEA2 results is less pronounced however, SPEA2 marginally achieves a better mean hypervolume than SPEA2-HDH at the cost of a high standard deviation. Statistical testing (Mann-Whitney U [31]) indicates the NSGAII and NSGAII-HDH results are significantly different, however there is no statistically significant difference between the final populations generated between SPEA2 and SPEA2-HDH. These results give a clear indication that the use of the HDH mutation tends to reduce variation in the performance of the algorithms, resulting in higher consistency.
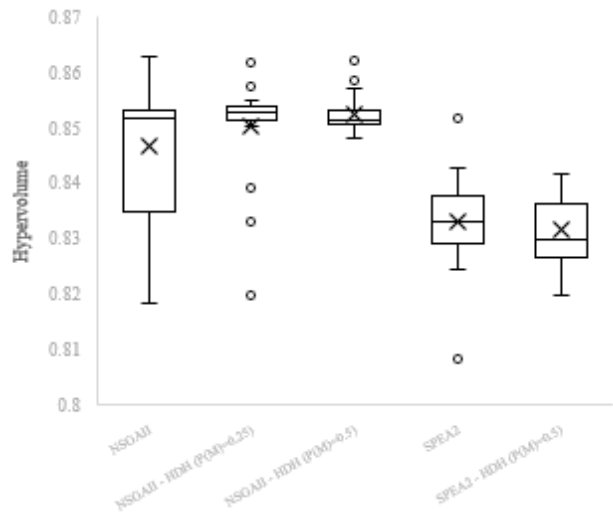

Figure 5: Hypervolume Results for the Hanoi Problem – NSGAII, NSGAII – HDH, SPEA2 & SPEA2 – HDH

The following set of experiments are for the Blacksburg problem, as with the Hanoi problem, each algorithm was run for 500,000 evaluations for 30 runs. Figure 6 shows the mean hypervolume of the algorithms over the runs. It can be observed that NSGAII-HDH (P(M)=0.25) exhibits slightly more aggressive convergence that its standard counterpart however the algorithms match performance after 250,000 evaluations. Interestingly when the HDH is applied to a greater extent (NSGAII-HDH (P(M)=0.5)) algorithm performance is diminished, resulting in slightly poorer performance. The performance of the SPEA2 based algorithms on the other hand differ greatly from the NSGAII based algorithms when compared to each other SPEA2 significantly outperform SPEA2-HDH throughout the search.
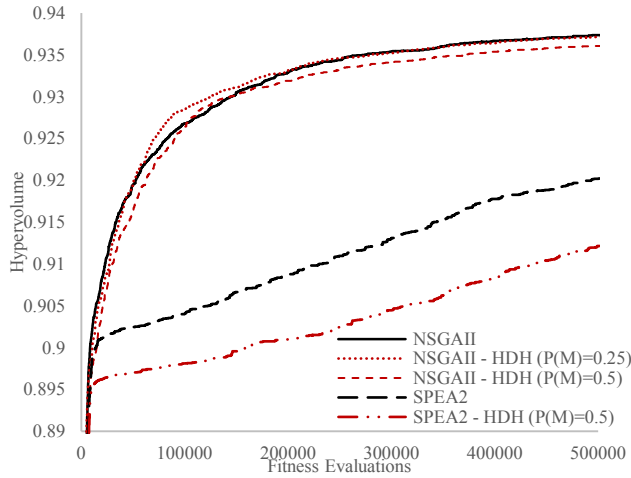
Figure 6: Mean Hypervolume for the Blacksburg Problem – NSGAII, NSGAII – HDH, SPEA2 & SPEA2 – HDH

Figure 7 illustrates very little distinction between the three NSGAII based algorithms, in fact, there is no statistically significant difference between the results. The difference between the SPEA2 based algorithms is more pronounced in this case, with the standard algorithm achieving a higher mean hypervolume and lower variance, a statistically significant result. The poor accuracy (31%) of the human derived model in this case is clearly having a detrimental effect on the performance of the algorithms, an effect more pronounced with SPEA2 than NSGAII.
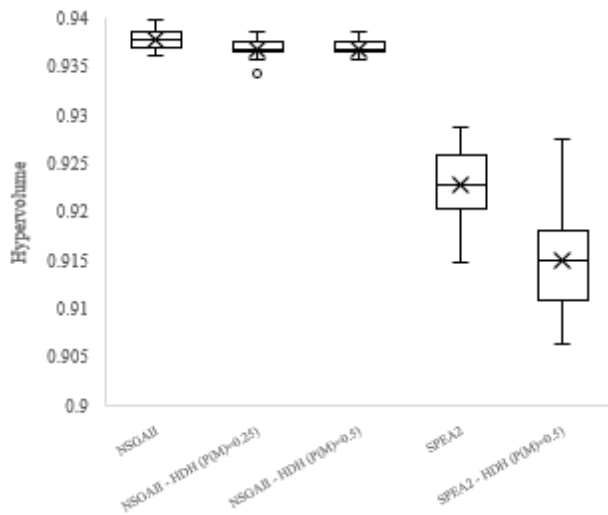

Figure 7: Hypervolume Results for the Blacksburg Problem – NSGAII, NSGAII – HDH, SPEA2 & SPEA2 - HDH

The final set of experiments were conducted on the Modena problem. In this case each algorithm was run 10 times for 500,000 fitness evaluations due to the greater computational cost of the hydraulic simulation for this network. Figure 8 shows the mean

hypervolume of the algorithms' populations over the allotted evaluations. From this figure NSGA-HDH (P(M)=0.5) can be seen to outperform NSGAII for the first 150,000 evaluations, displaying faster convergence. NSGAII then continues to gradually achieve a better hypervolume. However, NSGAII-HDH (P(M)=0.25) is shown to dominate NSGAII throughout the entirety of the search. SPEA2-HDH (P(M)=0.5) exhibits good performance in the early stages of the search compared with SPEA2, however SPEA2 surpasses the HDH variant at approximately 180,000 evaluations.
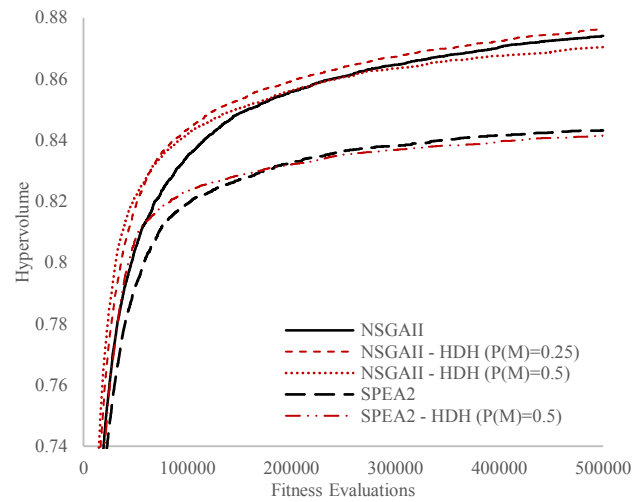

Figure 8 Mean Hypervolume for the Modena Problem – NSGAII, NSGAII – HDH, SPEA2 & SPEA2 – HDH

Figure 9 displays the final hypervolume results (10) for each of the algorithms. It can be seen that NSGAII-HDH (P(M)=0.25) achieves a better mean hypervolume with equal or lower variance compared with NSGAII and NSGAII-HDH (P(M)=0.5). Statistical analysis shows there is significant difference between NSGA-II-HDH (P(M)=0.5) and the other algorithms, however there is no significance between NSGAII and NSGAII-HDH (P(M)=0.25). SPEA2 achieves a slightly better mean hypervolume to NSHGAII-HDH (P(M)=0.5) with a lower variance.
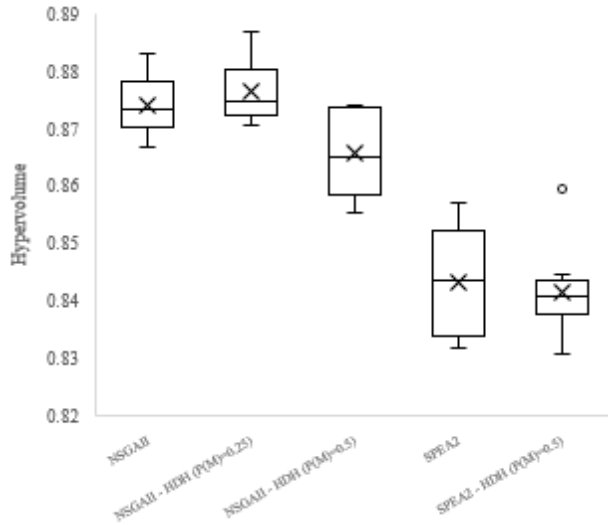
Figure 9: Hypervolume Results for the Modena Problem – NSGAII, NSGAII – HDH, SPEA2 & SPEA2 – HDH

From this set of experiments there seems to be a clear correlation between human derived model accuracy and the performance of the hybrid algorithms. The use of lower accuracy models, that are less capable of representing the human engineer with substantial accurately has been shown to have a detrimental effect on the search capabilities of an EA. It would appear that capturing heuristics from the larger models is more difficult where there are a greater number of potential inputs to consider. In this study, we have considered factors that are reasonably local to the selected pipe, however for larger networks the user may be making use of information from the extremities of the network outside of this locale. However even in the case of the Blacksburg and Modena problems there is no statistically significant performance decrease for NSGAII even though the models used in this study had an accuracy of below 50%. Furthermore, with careful parameter selection of HDH deployment, performance improvement on the large-scale Modena problem is possible throughout the optimization.

4   CONCLUSIONS

This paper demonstrates the potential for an EA to utilize human derived heuristics to improve performance when optimizing water distribution network design problems. The derived heuristics are automatically generated from the interactions of an engineer's manual alterations on a WDN and integrated into an EA through the mutation operator. The results show that the performance on an EA can be improved through the integration of domain specific expert knowledge. It is clear however that model accuracy is crucial to the effectiveness of the human-derived heuristics. As discussed, the human engineer is taking many features into account when making changes to a network, of which a number are not necessarily local to the pipe being manipulated. It is therefore important to take into account other network information when

developing new machine learning input schema. This however is a difficult task, not only in identifying the relevant information to feed the machine learning model but also to make it generalizable, so that a generic model can be trained and applied on multiple networks. Another consideration is which pipe to select given the current state of the network. It has been shown [32] that as the search of an EA progresses, the pipes close to the source(s) remain reasonably constant and the algorithm concentrates on the pipes closer to the extremities of the network. One of the key characteristics of a water distribution network is that the diameters of pipes close to the water source have a greater hydraulic influence over the whole system.

The findings presented in this paper describe an IEA that learns from its interactions with human experts, capturing knowledge and applying it to the process. With the development of a more generalizable input schema it is envisaged that knowledge gained on smaller networks like the ones presented in this paper could then be applied to very large WDN problems.

The developments presented in this paper open up the potential for more effective interaction between human expert and evolutionary algorithm resulting in better, more engineering feasible solutions to real-world problems. With further development the potential application of this approach could expand to a much larger and diverse set of problems beyond the field of water systems engineering.

ACKNOWLEDGMENTS

REFERENCES

[1]   M. B. Johns, E. Keedwell, and D. Savic, "Adaptive locally constrained genetic algorithm for least-cost water distribution network design," *J. Hydroinformatics*, vol. 16, no. 2, pp. 288–301, 2014.

[2]   K. McClymont, E. Keedwell, and D. Savic, "An analysis of the interface between evolutionary algorithm operators and problem features for water resources problems. A case study in water distribution network design," *Environ. Model. Softw.*, vol. 69, pp. 414–424, Jul. 2015.

[3]   W. Bi, G. C. Dandy, and H. R. Maier, "Improved genetic algorithm optimization of water distribution system design by incorporating domain knowledge," *Environ. Model. Softw.*, vol. 69, pp. 370–381, Jul. 2015.

[4]   C. Voudouris, E. P. Tsang, and A. Alsheddy, "Guided local search," in *Handbook of metaheuristics*, Springer, 2010, pp. 321–361.

[5]   C. Voudouris and E. Tsang, "Partial constraint satisfaction problems and guided local search," *Proc Pract. Appl. Constraint Technol. PACT96 Lond.*, pp. 337–356, 1996.

[6]   C. Voudouris and E. Tsang, "Guided local search and its application to the traveling salesman problem," *Eur. J. Oper. Res.*, vol. 113, no. 2, pp. 469–499, 1999.

[7]   P. Kilby, P. Prosser, and P. Shaw, "Guided local search for the vehicle routing problem with time windows," in *Meta-heuristics*, Springer, 1999, pp. 473–486.

[8]   E. Tsang and C. Voudouris, "Fast local search and guided local search and their application to British Telecom's workforce scheduling problem," *Oper Res Lett*, vol. 20, no. 3, pp. 119–127, 1997.

[9]   Q. Zhang, J. Sun, and E. Tsang, "An evolutionary algorithm with guided mutation for the maximum clique problem," *IEEE Trans. Evol. Comput.*, vol. 9, no. 2, pp. 192–200, 2005.

[10] S. Baluja, "Population-based incremental learning. a method for integrating genetic search based function optimization and competitive learning," Carnegie-Mellon Univ Pittsburgh Pa Dept Of Computer Science, 1994.

[11] H. Takagi, "Interactive evolutionary computation: Fusion of the capabilities of EC optimization and human evaluation," *Proc. IEEE*, vol. 89, no. 9, pp. 1275–1296, 2001.

[12] A. M. Brintrup, J. J. Ramsden, and H. Takagi, "Ergonomic chair design by fusing qualitative and quantitative criteria using interactive genetic algorithms," 2008.

[13] J. Nicklow *et al.*, "State of the art for genetic algorithms and beyond in water resources planning and management," *J. Water Resour. Plan. Manag.*, vol. 136, no. 4, pp. 412–432, 2009.

[14] A. Singh, B. S. Minsker, and A. J. Valocchi, "An interactive multi-objective optimization framework for groundwater inverse modeling," *Adv. Water Resour.*, vol. 31, no. 10, pp. 1269–1283, 2008.

[15] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, "A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II," in *International Conference on Parallel Problem Solving From Nature*, 2000, pp. 849–858.

[16] D. Winkler, M. Haltmeier, M. Kleidorfer, W. Rauch, and F. Tscheikner-Gratl, "Pipe failure modelling for water distribution networks using boosted decision trees," *Struct. Infrastruct. Eng.*, vol. 14, no. 10, pp. 1402–1411, Oct. 2018.

[17] J. Marques, M. Cunha, and D. A. Savić, "Decision Support for Optimal Design of Water Distribution Networks: A Real Options Approach," *Procedia Eng.*, vol. 70, pp. 1074–1083, Jan. 2014.

[18] Wang Qi, Guidolin Michele, Savic Dragan, and Kapelan Zoran, "Two-Objective Design of Benchmark Problems of a Water Distribution System via MOEAs: Towards the Best-Known Approximation of the True Pareto Front," *J. Water Resour. Plan. Manag.*, vol. 141, no. 3, p. 04014060, Mar. 2015.

[19] Prasad T. Devi and Park Nam-Sik, "Multiobjective Genetic Algorithms for Design of Water Distribution Networks," *J. Water Resour. Plan. Manag.*, vol. 130, no. 1, pp. 73–82, Jan. 2004.

[20] L. A. Rossman, "EPANET 2: users manual," 2000.

[21] O. Fujiwara and D. B. Khang, "A two-phase decomposition method for optimal design of looped water distribution networks," *Water Resour. Res.*, vol. 26, no. 4, pp. 539–549, 1990.

[22] H. D. Sherali, S. Subramanian, and G. V. Loganathan, "Effective Relaxations and Partitioning Schemes for Solving Water Distribution Network Design Problems to Global Optimality," *J. Glob. Optim.*, vol. 19, no. 1, pp. 1–26, Jan. 2001.

[23] C. Bragalli, C. D'Ambrosio, J. Lee, A. Lodi, and P. Toth, "On the optimal design of water distribution networks: a practical MINLP approach," *Optim. Eng.*, vol. 13, no. 2, pp. 219–246, 2012.

[24] M. B. Johns, D. J. Walker, E. Keedwell, and D. A. Savić, "Interactive Visualisation of Water Distribution Network Optimisation," presented at the International Conference on Hydroinformatics, 2018.

[25] D. J. Walker, M. B. Johns, E. Keedwell, and D. A. Savić, "Towards Interactive Evolution: A Distributed Optimiser for Multi- objective Water Distribution Network Design," presented at the International Conference on Hydroinformatics, 2018.

[26] D. J. Walker, "Generating Heuristics to Mimic Experts in Water Distribution Network Optimisation," *WDSA CCWI Jt. Conf. Proc.*, vol. 1, Jul. 2018.

[27] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, no. Oct, pp. 2825–2830, 2011.

[28] L. Breiman, J. Friedman, H, R. Olshen, A, and C. Stone, J, *Classification and Regression Trees*. Routledge, 1984.

[29] R. Farmani, D. A. Savic, and G. A. Walters, "Evolutionary multi-objective optimization in water distribution network design," *Eng. Optim.*, vol. 37, no. 2, pp. 167–183, 2005.

[30] E. Zitzler and L. Thiele, "Multiobjective optimization using evolutionary algorithms—a comparative case study," in *International conference on parallel problem solving from nature*, 1998, pp. 292–301.

[31] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *Ann. Math. Stat.*, pp. 50–60, 1947.

[32] E. Keedwell, M. Johns, and D. Savic, "Spatial and temporal visualisation of evolutionary algorithm decisions in water distribution network optimisation," in *Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation*, 2015, pp. 941–948.