Integrated Computer-Aided Engineering 19 (2012) 111–124
DOI 10.3233/ICA-2012-0395
IOS Press

Causally-guided evolutionary optimization and its application to antenna array design

Timur Chabuk[a,*], James Reggia[a], Jason Lohn[b] and Derek Linden[b]
[a]*Department of Computer Science, University of Maryland, College Park, MD, USA*
[b]*Carnegie Mellon University, NASA Research Park, Mountain View, CA, USA*

Abstract. In recent years, evolutionary computation has been successfully used to solve problems involving engineering design and invention, sometimes producing results that are qualitatively different than previous traditionally-designed solutions. However, while evolutionary methods appear to be a promising tool for supporting design, their usefulness is substantially limited by their computational expense and inability to integrate expert knowledge with evolutionary search. Here we develop and evaluate methods for causally-guided evolutionary design based on expert-supplied cause-effect relations that guide how genetic operators are applied (in contrast to conventional genetic operations which are carried out blindly and randomly), using these methods for antenna array design. To our knowledge, this is the first study that biases genetic operations in response to the specific performance characteristics of the individuals to which they are applied, and the first to use explicit cause-effect relations to guide this process. Our experimental evaluation compares using evolutionary systems with and without causal guidance to design directional dipole antenna arrays that meet pre-specified performance criteria. We find that causally-guided systems produce optimal solutions with significantly greater frequency and significant computational savings, suggesting that this approach may substantially improve the use of evolutionary computation in engineering design.

Keywords: Causal knowledge, antenna array design, knowledge-guided evolution, causally-guided evolution

1. Introduction

During recent years, there has been much interest in using computational intelligence optimization methods for engineering design, including evolutionary computation, particle swarm optimization [12,41], ant colony optimization [38], and other nature-inspired approaches. Evolutionary computation in particular has been successfully used to design electronic circuits [27,46], antennas [5,31–34], space truss structures [2,4], control mechanisms for unmanned aerial vehicles [43], wind turbines [29], high-rise building structures [3], and nuclear power plant monitoring plans [7]. The design process in these situations is genuine human-machine collaboration: a person defines the problem, search space, fitness function, and other design constraints, while the evolutionary process generates and evaluates a much larger number of alternative designs than could be done manually. Sometimes the results of evolutionary systems are even qualitatively different from previous human-only solutions, such as patentable electronic circuits [27], and novel irregularly shaped antennas [21].

Evolutionary computation thus appears to be a very promising tool for supporting the engineering design process. However, in order for the evolutionary process to remain computationally tractable when applied to increasingly complex engineering design problems, new extensions must be developed that increase the efficiency and effectiveness with which evolutionary systems produce optimal designs. To this end, the goal of the research presented here is to develop one such potential extension, "causally-guided evolution", and to evaluate its potential usefulness by applying it to a real-world antenna array design problem.

*Corresponding author: Timur Chabuk, Department of Computer Science, University of Maryland, College Park, MD, USA. E-mail: Chabuk@cs.umd.edu.

ISSN 1069-2509/12/$27.50

By causally-guided evolution, we mean an evolutionary system in which genetic operators use expert-supplied causal knowledge to bias, but not to control, the modification of individual designs. Prior to the beginning of the evolutionary process, human experts supply knowledge in the form of cause-effect relations from the domain. Each cause-effect relation represents a rule about how some part of an individual's genetic representation relates to some part of the same individual's phenotypic performance. During the evolutionary process, when a genetic operation is applied to an individual, the performance characteristics of the individual are examined and the cause-effect relations are used to identify likely design problems in the genotype of the individual. The genetic operation is then probabilistically biased toward modifying the individual's specific design problems, For example, in applying a mutation operator to an individual parent to produce a modified offspring, the mutation will be biased so that those parts of the genotype that are judged more likely to be flawed are made more likely to be mutated. Similarly, in applying a crossover operator to two parents, the operation is biased so that those parts of the parents that are judged to have less chance of being flawed are made more likely to be combined together in the offspring.

This is the key distinction between the work presented here and past studies of evolutionary systems; our genetic operators examine the performance characteristics of each individual to which they are applied, and are biased based on the specific design problems that are reasoned to be present in each individual. In contrast, the genetic operators in past evolutionary studies are applied without any regard to the performance characteristics of individual designs. In traditional evolutionary methods the genetic operators are blind and random, whereas here they are guided. While genetic operators in past studies on "knowledge incorporation" have been specially designed based on expert knowledge, when applied to individuals the operators still execute without regard to the performance characteristics of the individuals [8,13,17,25,44,45].

Our hypothesis is that causally-guided genetic operators as described above will ultimately make evolutionary systems more effective by allowing them to explore a much larger number of good designs while still exploring novel solutions that initially appear unpromising, and also more computationally efficient by decreasing the number of poorly fit individuals that do not contribute useful information to the evolutionary process. We believe that the benefits of causally-guided evolutionary systems will be most pronounced when applied to problems in which domain expertise is present but insufficient for solving problems in closed form. Causally-guided evolutionary methods are designed to preserve the limited dependence on domain knowledge found in traditional evolutionary computation, while leveraging whatever cause-effect knowledge is available.

Of course, there is no guarantee that causal guidance will improve the effectiveness or efficiency of evolutionary methods, and it is entirely possible that just the opposite would be true; adding causal influences could produce evolutionary search that is less effective, less creative, and less computationally efficient. Further, incorporating expert knowledge into an evolutionary system raises the possibility of producing non-optimal solutions [6]. Of particular concern here is the possibility that causal-guidance will steer the evolutionary process into local optima, preventing the discovery of globally optimal solutions. To assess this issue, we compare causally-guided versus non-causally guided evolutionary design of an antenna array in the following.

2. Methods

2.1. Causally-guided evolutionary computation

Causally-guided evolution uses causal knowledge supplied by domain experts prior to the beginning of the evolutionary process. Each piece of knowledge details a cause-effect relationship, expressed as

$$\text{Genotypic Disorder} \rightarrow \text{Phenotypic Symptom}$$

The arrow here is not logical implication, but causality. A *genotypic disorder* is simply a non-optimal part of an individual's genetic material, while a *phenotypic symptom* is a particular performance or fitness problem that an individual may have. The term "phenotype" is used broadly here, to include not only the structure/form that the genotype develops into, but also the behavioral/performance characteristics of that form. Domain experts do not need to supply an exhaustive list of cause-effect relations, and could even supply as little as a single relation; they only need to supply those causal relationships of which they are aware and believe to be most important.

In causally-guided evolution, once it has been determined in the usual fashion that a genetic operator will be applied to a specific individual, the individual's per-

Generic Causally-Guided Genetic Operator
Step 1: Examine the individual's performance characteristics to identify its phenotypic symptoms/flaws.
Step 2: Based on supplied causal relations and the individual's phenotypic symptoms, determine what parts of the individual's genome are likely to be flawed.
Step 3: Bias the application of the genetic operator to the individual accordingly

Fig. 1. High level overview of the three step process that is followed by causally-guided genetic operators.

formance characteristics and the expert-supplied causal knowledge are used to bias how the operator is applied to the individual. This is accomplished in three steps, as shown in Fig. 1. The exact manner in which the genetic operations are biased in step 3 depends on the specific type of the genetic operator in question, as follows.

Causally-guided mutation operations are biased so that those parts of the genotype with higher relative likelihoods of being flawed are made more likely to be mutated. In this way, causally-guided mutation increases the chances that problematic genes will be changed. Causal guidance does not change the number of modifications that will be made to an individual's genotype, only where they are made.

Causally-guided crossover operations are biased so that those parts of the parent individuals' genotypes that have lower relative likelihoods of being flawed are made more likely to be combined together when creating offspring. Consequently, those parts of the individuals' genotype that have higher relative likelihoods of being flawed are made less likely to be used when creating offspring. Causally-guided crossover increases the chances that the best parts of parents are combined into the produced offspring.

Causal guidance is used to bias genetic operations, but does not explicitly control them. Causal guidance is applied probabilistically and does not prevent the occurrence of poorly fit individuals that arise in the populations due to random alterations; it simply influences the process towards the formation of more fit individuals and fewer very poor individuals than would otherwise occur.

To assess the effectiveness of these ideas, we explore the use of causally-guided evolution to solve an antenna array design problem. A generational genetic algorithm augmented with causally-guided genetic operators was developed, and its performance in solving the antenna array design problem was compared to a carefully-matched genetic algorithm that uses no causal-guidance but is otherwise equivalent. With the exception of the causally-guided genetic operators, all aspects of the evolutionary systems (described in more detail below) are conventional and widely used [13]. The goal of these experiments is to determine if causally-guided genetic operators mislead the evolutionary process toward local minima, have no significant effect, or improve the quality of and speed with which solutions are produced.

2.2. Antenna design

The design of many real-world antennas and antenna arrays is difficult because it requires significant domain expertise and it is time and labor intensive [14, 16,42]. Complex interactions between neighboring components of an antenna can make it very difficult to predict antenna behavior beforehand. Past work has explored using particle swarm optimization [40], ant colony optimization [38], simulated annealing [12] and other automated methods for designing antennas. Evolutionary computation methods in particular have been successfully used to design Yagi-Uda antennas, quadrifiliar antennas, and crooked wire monopole antennas [19,31–34]. To our knowledge, no antenna has ever been designed through causally-guided evolution, as described here.

There are a number of performance characteristics that are important in antenna design. The term *directivity* refers to the capability of an antenna to radiate more energy in certain directions than in others. *Gain* is a measure of the amount of energy that an antenna radiates in a specific direction. It is calculated by computing the ratio of energy radiated in that direction to the amount that would be radiated by an antenna that radiates equally in all directions. Gain often has very high values and is most often expressed in decibels (dB). Another important characteristic is the impedance mismatch between transmission lines and antenna, which can cause electrical signals to reflect back through the feed network. *Voltage standing wave ratio* (VSWR) quantifies impedance mismatch between transmission lines and radiating elements (VSWR $= 1$ is ideal). Finally, with antenna design problems, *cost* can mean many different things, including manufacturing difficulty, weight, size, volume of material, etc.

2.3. Dipole antenna array

The specific task used in this work is that of designing a directional dipole antenna array that meets prespecified performance criteria. Dipole antenna arrays consist of an array of parallel lengths of wires, known as dipoles, which are positioned above a ground plane. A transmission line connects to the center of each dipole and carries the signal that is radiated or received by the antenna. The complete design specifications for such an antenna includes the number of dipoles, the lengths of dipoles, the height of dipoles off the ground plane, the spacing between dipoles, and the phases and voltages with which each dipole is fed. For this work, dipole antenna arrays were limited to having uniform spacing, height, and length. Known formulas are used to calculate the desired voltage and phase with which each dipole should be fed, based on each dipole's location and the desired direction of broadcast. The uniform nature of such designs makes them appear to be very simple. However, in practice and despite the small dimensionality of the search space, it is quite difficult for a human designer to optimize these four values by hand. Furthermore, greedy or local search algorithms often get stuck in the many local optima that exist.

In the particular antenna array design task considered here, the specific goal, provided a-priori, is to maximize gain between -10 and $+10$ degrees off boresight in the plane that bisects the dipoles, minimize VSWR, and minimize cost. Specifically, a successful antenna must have an average gain of at least 10 dB in the target angle range and a VSWR of less than 3.0 (a commonly used limit for VSWR in antenna design). The number of dipoles in the antenna array is used as a rough approximation for cost, which should be minimized but does not have a required value. The antenna is to be operated at 1200 MHz with 50 ohm transmission lines. These particular performance requirements were selected because they define an antenna design problem that is complex enough to be of real-world interest while remaining simple enough for an initial exploration into causally-guided evolutionary computation. As in previous studies, the antennas are simulated over an infinite ground plane in order to keep computational costs down [34]. Software was implemented in Java and C, and runs on Linux-based PC's. All antennas were simulated using an open source version of the Numerical Electromagnetic Code software package [10].

2.4. *Fitness function and genetic representation*

A fitness function that captures the performance criteria above, is made up of three components:

$$F_{Overall} = F_{VSWR} + F_{Directivity} + F_{Cost} \quad (1)$$

The VSWR component rewards low VSWR values, and was calculated as:

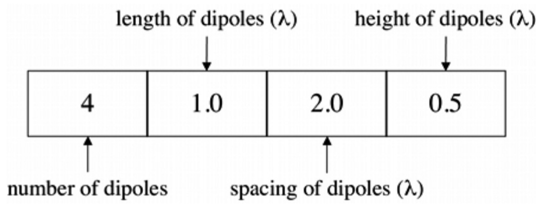

Fig. 2. Genetic representation of a dipole antenna array.

$$F_{VSWR} = \begin{cases} -2 * \mathrm{VSWR}_{\max} & \text{if } \mathrm{VSWR}_{\max} \geqslant 3.0 \\ -1 * \mathrm{VSWR}_{\max} & \text{if } \mathrm{VSWR}_{\max} < 3.0 \end{cases} \quad (2)$$

where $\mathrm{VSWR}_{\max}$ is equal to the maximum VSWR observed at any dipole in the antenna. If the $\mathrm{VSWR}_{\max}$ is above 3.0, the F_{VSWR} score is multiplied by -2 instead of -1, increasing its negative impact on the overall fitness value. The $\mathrm{F}_{Directivity}$ score measures the directivity of the antenna in the target angle range, and is calculated as:

$$F_{Directivity} = \sum_{i=-10}^{10} Gain_i \quad (3)$$

where Gain_i is the amount of gain observed in the XY-plane at i degrees off boresight. There are 21 terms here, each in the -10 to $+10$ degree range, so an antenna that meets the design requirements of 10 dB average in this range would have a $\mathrm{F}_{Directivity}$ score of at least 210. Lastly, the F_{Cost} component is equal to the number of dipoles in the antenna array multiplied by -1. These three components are summed to yield the overall fitness score ($\mathrm{F}_{Overall}$), which is maximized by the evolutionary system. This fitness function and the stepped nature of the VSWR component are inspired by previous studies showing such an approach to be successful [21,34].

The genetic representation used to represent a dipole antenna array consists of a vector of four numbers (see Fig. 2). The first is a whole number between 1 and 10 that represents the number of dipoles. The remaining real-valued numbers may have values between 0.1 and 2.0. These three values represent distances, expressed in wavelengths, for the length of dipoles, height of dipoles and spacing between dipoles. Since the operating frequency of this antenna is 1200 MHz, a wavelength corresponds to roughly 0.25 m. For example, an individual in the population with a genetic vector of [4, 1.0, 2.0, 0.5] describes an antenna with 4 dipoles of 0.25 m length, spaced 0.5 m apart and 0.125 m off the ground plane. Each vector element is a "gene". The letters N, L, S, and H are used to refer to these four genes.

```
function CAUSAL-MUTATION(Individual A)
  Let G = {N,L,H,S}
  // STEP 1: assess phenotypic symptoms
  if (A has VSWR > 3.0) then
      hasSymptom(A, High VSWR) = true
  else
      hasSymptom(A, High VSWR) = false
  end if
  // STEP 2: infer likelihood genotype disorders
  increasedLikelihood(A,d(g))=false,∀g∈G
  decreasedLikelihood(A,d(g))=false,∀g∈G
  if (hasSymptom(A, High VSWR) then
      increasedLikelihood(A,d(L)) = true
  else
      decreasedLikelihood(A,d(L)) = true
  end if
  // STEP 3: bias mutation
  utility(s) = 1.0 ∀ specific mutations s
  for all genes g ∈ G do
      if(increasedLikelihood(A,d(g))) then
          utility(M_g) = utility(M_g) * Δ1
      else if (decreasedLikelihood(A,d(g))) then
          utility(M_g) = utility(M_g) / Δ2
      end if
  end for
  rescale all utility scores to sum to 1.0
  for all specific mutations s do
      if RANDOM(0,1) < utility(s) then
          apply s to A
      end if
  end for
```

Fig. 3. Pseudocode for the causally-guided mutation operator as implemented for the dipole antenna array design task. Input argument A is an individual antenna array design, while d(g) represents the assertion that gene g of A is flawed (see text). Mutation s refers to the 8 specific mutations listed in this text. Constants Δ_1 and Δ_2 are constrained to be greater than 1.0, and RANDOM(0,1) returns a uniformly random floating point number in (0.0, 1.0).

2.5. *Cause-effect relations*

The non-linear interactions between antenna features and performance qualities are not easily characterized. Here the simple case in which only a few pieces of causal knowledge are incorporated into the system is considered. This allows an initial exploration into the feasibility of causally-guided evolution, while delaying the need to seriously address complex situations in which multiple genotypic disorders influence multiple phenotypic symptoms simultaneously. It is known by antenna design experts that sub-optimal dipole lengths can cause high VSWR, but that the height, length, and number of dipoles have limited effect on the VSWR of an antenna. This is the causal knowledge that we use in this study:

Sub-optimal Dipole Length → High VSWR

The left side of this causal relation is a genotypic disorder because it refers to a sub-optimal gene value in the genetic representation (Fig. 2). As noted earlier, the → symbol is not logical implication but causality. There are three additional genotypic disorders: sub-optimal number of dipoles, sub-optimal spacing of dipoles and sub-optimal height of dipoles. These three genotypic disorders do not cause high VSWR. The terms d(N), d(L), d(S) and d(H) are used to represent the four genotypic disorders.

2.6. *Causally-guided genetic operators*

As outlined in Fig. 1, causally-guided genetic operations occur in three steps. In the antenna array design task, assessing phenotypic symptoms (Step 1) and making inferences about the likelihood of genotypic disorders (Step 2) is straightforward. If an antenna has a VSWR greater than 3.0, it is assessed as having the phenotypic symptom of high VSWR, otherwise it is assessed as *not* having the symptom of high VSWR. An antenna that has the symptom of high VSWR can be reasoned to have an increased chance of having sub-optimal dipole lengths, as this is the only genotypic disorder that is known to cause the symptom. On the other hand, if the antenna lacks the symptom of high VSWR, it can be reasoned that there is a decreased chance that the antenna has sub-optimal dipole lengths. Finally, the way in which these inferences are used to bias the execution of genetic operators (Step 3) depends on the particular genetic operator in question.

2.6.1. *Causally-guided mutation*

When an individual is selected for causally-guided mutation, there are a number of specific mutations that may or may not be applied to the individual, and these are controlled by the algorithm Causal-Mutation in Fig. 3. The decision of whether to apply each specific mutation is made stochastically and independently, and is influenced by causal guidance. There are eight specific mutations that may be applied to an individual during each causally-guided mutation: for each of the four genes there is a specific mutation that makes large changes to the gene value, and one that makes small changes. The terms M_N, m_N, M_L, m_L, M_S, m_S, M_H, and m_H are used to refer to these eight specific mutations. A lowercase m is used for small mutations, an uppercase M is used for large mutations, and single character indices are used to identify the relevant gene.

Each large mutation replaces the relevant gene with a random value, selected uniformly from the appropriate legal range for that gene. Thus large specific mutations very often (though not always) make large changes to the gene value. If m_N is applied to an individual, the value of the gene that specifies the number of dipoles is either incremented or decremented by 1, with equal probability. For the other three genes, small mutations either increase or decrease the gene's value by a small random value that is selected with equal probability from one of five uniform distributions: [0, 0.2], [0, 0.1], [0, 0.01], [0, 0.001], and [0, 0.0001]; this has the effect of making smaller changes (e.g., < 0.0001) more probable than other small changes (e.g., 0.1 to 0.2).

When causally-guided mutation is applied to an individual, the probability with which each of these specific mutations is applied is biased based on the inferred likelihood of the various genotypic disorders (see Fig. 3). Specifically, a utility score is calculated for each specific mutation. The utility score of a specific mutation is used as an indication of how useful it would be to apply that specific mutation to the individual. Initially, each specific mutation is assigned a utility score of 1.0. For each gene that has an increased likelihood of being flawed (i.e., the corresponding genotypic disorder has an increased likelihood of being present), the utility score of the corresponding large mutation is increased by multiplying it by a constant $\Delta_1 > 1.0$. For those genes with lower likelihoods of being flawed, the utility score of the corresponding large mutation is decreased by dividing it by a constant $\Delta_2 > 1.0$. Lastly, the utility scores are rescaled so that the sum of all eight specific mutations' utility scores is equal to one. Thus large mutations that correspond to genes with higher relative likelihoods of being flawed have higher utility scores. Each specific mutation is then applied with probability equal to its utility score. For example, when causally-guided mutation is applied to an antenna design with the symptom of high VSWR, the causal knowledge indicates an increased probability that dipole lengths are sub-optimal. Thus, the utility score of the specific mutation M_L, which makes changes to the gene associated with dipole length, is increased, while the utility scores of all other specific mutations are effectively decreased through the normalization process.

2.6.2. *Causally-guided crossover*

As implemented in algorithm Causal-Crossover (see Fig. 4), causally-guided crossover is a variation of uniform crossover, in that offspring are created by stochastically selecting one copy of each gene from each of two parents M and F, and each gene is inherited independently. The inferred likelihood of the genotypic disorders is used to bias the crossover operation as follows. Each gene in each parent is initially assigned an *inheritance score* (inh) of 1.0. This score is designed to be an indication of how *un*likely it is that the gene is sub-optimal and, accordingly, how useful it would be for offspring to inherit that gene. The inheritance score of each gene that has an increased likelihood of being flawed is decreased by subtracting the constant $0 < \Delta_3 < 1$. This same constant is added to the inheritance score of each gene that has a decreased likelihood of being flawed. The chance that an offspring will in-

```
function CAUSAL-CROSSOVER(Individual M, F)
    let P = {M, F}, G = {N, L, H, S}
    // STEP 1: assess phenotypic symptoms
    for all p ∈ P do
        if (p has VSWR > 3.0) then
            hasSymptom(p, High VSWR) = true
        else
            hasSymptom(p, High VSWR) = false
        end if
    end for
    // STEP 2: infer genotype disorders likelihood
    increasedLikelihood(p,d(g))=false∀g∈G,∀p∈P
    decreasedLikelihood(p,d(g))=false∀g∈G,∀p∈P
    for all p ∈ P do
        if(hasSymptom(p, High VSWR) then
            increasedLikelihood(p,d(L)) = true
        else
            decreasedLikelihood(p,d(L)) = true
        end if
    end for
    // STEP 3: bias crossover
    create child Individual C
    inh(p,g) = 1.0 ∀g∈G, ∀p∈P
    for all p ∈ P do
        for all genes g ∈ G do
            if increasedLikelihood(p,d(g)) then
                inh(p,g) = inh(p,g) – Δ3
            else if decreasedLikelihood(p,d(g)) then
                inh(p,g) = inh(p,g) + Δ3
            end if
        end for
    end for
    for all g ∈ G do
        if RANDOM(0,1) < inh(M,g) / (inh(M,g) + inh(F,g)) then
            replace C's value of g with M's
        else
            replace C's value of g with F's
        end if
    end for
```

Fig. 4. Pseudocode for the causally-guided crossover operator as implemented for the dipole antenna array design task. Same notation as in Fig. 3, where $0 < \Delta_3 < 1$ is a constant. A parent p's inheritance score for gene g is given by inh(p,g), as explained in text.

herit a copy of a gene from one parent is equal to that parent's gene's inheritance score divided by the sum of both parents' genes' inheritance scores. In this manner, those genes in a parent with high relative likelihoods of being flawed will have lower inheritance scores and therefore be less likely to be inherited by offspring. For the antenna array design task, the chromosome consists of only four genes, which limits the potential for overly destructive crossover.

2.6.3. *Control genetic operators*

The control mutation and crossover operators are not biased by causal reasoning, but otherwise operate the same as their causally-guided counterparts. They serve as a baseline for comparing the effectiveness of causally-guided operators. During *control mutation*, each of the same eight specific mutations may or may not be applied to an individual with fixed probability of 1/8. During *control crossover*, offspring are created by stochastically selecting one copy of each gene from one of the two parents, with each gene having a fixed, equal chance of coming from either of the two parents.

2.7. *Experimental methods*

Four different evolutionary systems were applied to the dipole antenna array design problem. The CONTROL system used control crossover and control mutation operators; it makes no use of causally-guided genetic operators. In contrast, causal mutation and control crossover were used by the $\mathrm{CAUSAL_M}$ system, control mutation and causal crossover were used by the $\mathrm{CAUSAL_C}$ system, and causal mutation and causal crossover were used by the $\mathrm{CAUSAL_{CM}}$ system. Two hundred trials of each of these four evolutionary systems were conducted using a different random number stream for each trial. Each trial was started with a randomly generated population and executed for 1000 generations, yielding 50,000 antenna array simulations. Individuals in the initial population were created by selecting gene values uniformly from the range of legal values. A population of size 50 and tournament selection with tournament size two were used in all trials. In each generation, exactly one offspring was created by elitism. Each of the remaining 49 offspring in each generation was created by using exactly one of the following stochastically chosen operators: crossover (47.5% chance), mutation (47.5%), or reproduction (5%). Thus, the number of offspring created by each method in each generation was not constant or predetermined. However, given a population size of 50, in each generation the expected number of individuals created by crossover, mutation and reproduction were approximately 23.3, 23.3, and 2.4, respectively (with an additional single offspring via elitism). The constant values Δ_1, Δ_2, and Δ_3 were fixed at 20.0, 2.0, and 0.2, respectively. These parameter values were found via a small number of test runs; they may not be optimal, but were found to be effective in this study, and were the same in both control and experiment trials.

The results of the 800 trials (200 trials times 4 evolutionary systems) were examined and analysis was performed as follows. Of all the antenna designs produced by the evolutionary systems, there appears to be a clear delineation between those that reach a fitness level of 310 and those that do not (i.e., the fitness of the best antenna produced by each evolutionary process is either just above 310 or else is considerably lower. Thus, a fitness level of 310 offers useful criteria by which to classify antennas as being "optimal" for the purpose of analysis. Data was collected to determine how often each of the four evolutionary systems was able to find an optimal antenna within various numbers of generations. Additionally, the average number of generations required by each system to find an optimal antenna was calculated. When computing these averages, trials that did not find an optimal antenna design in 1000 generations were counted as having found one in generation 1000. Therefore, this average is actually a rough approximation of the true average.

Multi-start strategies, in which evolutionary processes are terminated and restarted if an adequate solution is not found within a certain generational limit, often find adequate solutions faster than by running a single process indefinitely [21]. A multi-start strategy was not used in the 200 trials of each system performed in this work. However, by using the results of each system's 200 trials as an approximation for how that system performs across all random number streams, one can calculate the expected number of generations required by each system to find an optimal antenna when used in conjunction with a multi-start strategy, thus providing a measure of computational cost that is very practical. To do this, we calculated E(required_gens(S,f,g)), the expected value of the number of generations required by system S to find an antenna with fitness f when used with a multi-start strategy and a generation limit of g as in Eq. (4). Here success_rate(S,f,g) is the fraction of system S's 200 trials that find an antenna with fitness of at least f by generation g, and average_gens(S,f,g) represents the average number of generations required by these trials.

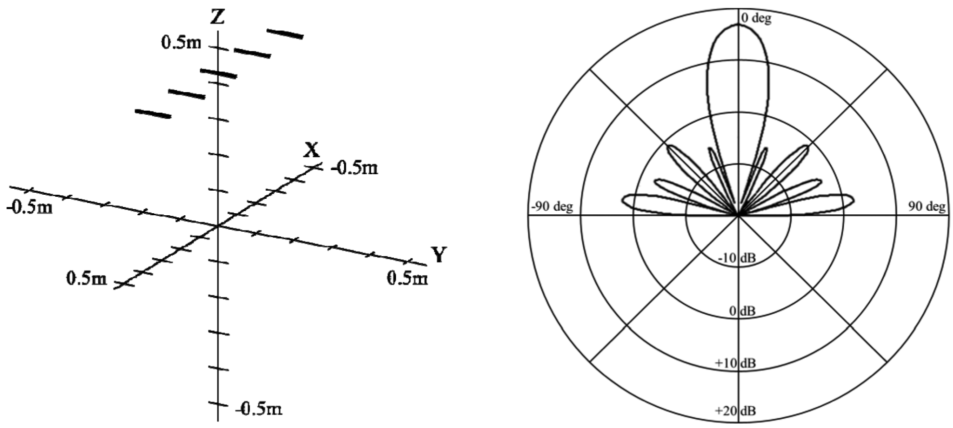

Fig. 5. A schematic of the fittest evolved antenna (left) and a radiation plot (right) illustrating the antenna's directivity. The ground plane is located in the XY plane (not shown).

$$\text{E(required-gens(S,f,g))} = \frac{1}{\text{success-rate(S,f,g)}} - g + \text{average-gens(S,f,g)} \quad (4)$$

Equation (4) follows from probability formulas related to Bernoulli trials, as each evolutionary process can be thought of as a Bernoulli random variable and the multi-start approach is a Bernoulli trial. Because it is difficult to know a good generation-limit value *a priori*, the expected number of generations required by each system to find an optimal antenna was calculated with a variety of generation-limits: 100, 200, 300, 400 and 500. This gives a very practical measure of the different computational costs associated with using each of the four evolutionary systems to find an optimal antenna design.

The various antenna designs produced by the four systems were examined. The antennas were visually inspected, found to fall into clusters according to their genotype similarity, and the clusters were assigned arbitrary labels. The frequencies with which the various systems arrived at these different designs were calculated, in an effort to understand the ways in which causal-guidance affects the types of designs that are produced.

3. Results

Numerous trials from all four evolutionary systems successfully designed antennas that met the prespecified performance criteria. The most fit individual, which was discovered by some trials in each system, was a five-element dipole array with dipoles of length 0.4635 λ, height of 1.7094 λ off the ground plane, and spacing of 0.6956 λ. This antenna had a VSWR of only 1.31 and a total directivity score of just over 316, indicating an average of just over 15 dB of gain in the target range. A schematic of this antenna design and a radiation plot illustrating its directivity can be seen in Fig. 5. The ground plane, which would occupy the XY-plane where Z = 0, and the feed-lines are not pictured. The radiation plot can be thought of as corresponding to the XZ-plane where Y is equal to 0, which bisects the dipoles (main lobe points in the positive Z-direction). The radiation plot is in terms of gain, which simply shows relative strength in particular directions. The overall fitness score of this most fit antenna was 310.04, compared to typical fitness values of 250 to 290 in the initial generation. The distribution of the fitness values of the evolved antennas is such that there is a clear delineation between the fittest antennas and the less fit ones. For the most part, evolved antennas either have a fitness score of just over 310 or a fitness score that is much lower ($<$ 308.5). As noted earlier, any antenna with a fitness of 310 or higher is considered to be an "optimal" antenna design.

A higher percentage of the causally-guided evolutionary systems' trials than control systems' trials found an optimal antenna design within 1000 generations. An individual trial is said to be successful by generation g if it finds an optimal antenna design (as defined above) at or before generation g. Figure 6 illustrates, for each of the four evolutionary systems, the fraction of trials that were successful by generation 100, 250, 500 and 1000. At each generation listed, the causally-guided systems found optimal antennas with greater frequency than the control system. Further, the perfor-

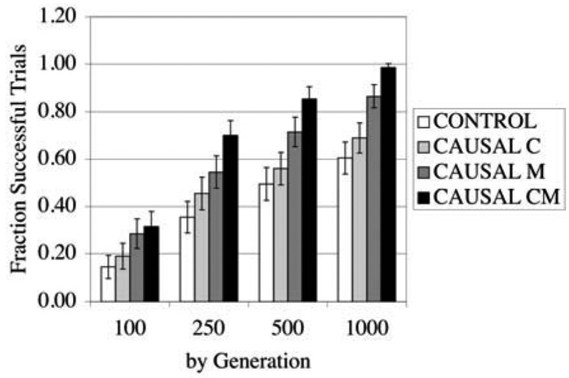

Fig. 6. Fraction of trials of each system that find an optimal antenna design within 100, 250, 500 and 1000 generations. Vertical bars are used to illustrate a 99% confidence level.

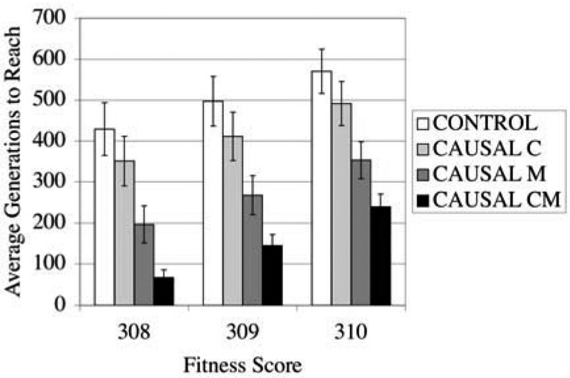

Fig. 7. The average number of generations that each system requires to find antenna designs of various fitness scores. Vertical bars are used to illustrate a 99% confidence interval.

mance of the four systems relative to each other appears to be the same in all generations. $\mathrm{CAUSAL_{CM}}$ outperforms $\mathrm{CAUSAL_{M}}$, which outperforms $\mathrm{CAUSAL_{C}}$, which outperforms the CONTROL system. A z-test revealed the difference between $\mathrm{CAUSAL_{CM}}$ and CONTROL to be statistically significant at a 99% confidence level at generation 250, 500 and 1000. At generation 1000 and 500, the difference between all pairs of systems were statistically significant at a 95% confidence level, except for CONTROL and $\mathrm{CAUSAL_{C}}$, which still had a low p-value of less than 0.10 in generation 1000. Note that the 99% confidence intervals of CONTROL and $\mathrm{CAUSAL_{CM}}$ never overlap.

The average number of generations required by each evolutionary system to find antenna designs with scores of 308, 309, and 310 are illustrated in Fig. 7. For each fitness score, the $\mathrm{CAUSAL_{CM}}$ system averaged the lowest number of generations, followed by $\mathrm{CAUSAL_{M}}$, $\mathrm{CAUSAL_{C}}$, and CONTROL. The $\mathrm{CAUSAL_{CM}}$ system averaged less than 16%, 29%, and 42% as many generations as CONTROL. The difference between $\mathrm{CAUSAL_{CM}}$ and all other systems, as well as $\mathrm{CAUSAL_{M}}$ and all other systems, was statistically significant to a 99% confidence level, for each fitness score. The difference between the $\mathrm{CAUSAL_{C}}$ system and the CONTROL system was statistically significant to a 95% confidence level.

By using the 200 trials for each system as an approximation for how the system performs across all initial random seeds, one is able to calculate the expected number of generations required to find an optimal antenna, when used in conjunction with multi-start strategies using generation limits of 100, 200, 300, 400 and 500 (see Section 2.7). As illustrated in Fig. 8, all of the causally-guided systems outperformed the CONTROL system, regardless of the generation limit was used. The $\mathrm{CAUSAL_{CM}}$ system has the lowest expected value, followed by $\mathrm{CAUSAL_{M}}$ and $\mathrm{CAUSAL_{C}}$. The $\mathrm{CAUSAL_{CM}}$ system requires less than 43% as many generations as the CONTROL system, regardless of generation limit.

It was also found that each of the most fit antenna designs that were produced by each trial of each evolutionary system may be grouped, based on the similarity of their genotypes, into one of seven design categories. The mean values of the design aspects and performance characteristics of antennas from each of these categories are detailed in Table 1, and the categories are assigned arbitrary labels A through G. In categories A through F, there is very little variation among antenna designs. The maximum Euclidean distance of any antenna from the average characteristics of its assigned category is less than 0.08. Category G captures 3 outlier antennas that do not fit into any of the other six categories. These outlier antennas are tightly clustered in terms of dipole length, dipole height, and dipole spacing but, unlike antennas from the other six categories, may have different numbers of dipoles (9 or 10). Category-A antennas represent the fittest class of antenna designs. The other categories of antennas represent local optima at which the evolutionary systems sometimes got stuck.

There appears to be little difference between the types of antenna designs that were evolved by the control system and causal systems. With the exception of category-G antennas, which are only 3 out of 800 evolved antennas, there are no antenna designs that were produced by the causal systems that were not evolved, in at least one trial, by the control system. However, there are significant differences in the fre-

Table 1
The seven classes of antenna designs produced by all evolutionary systems

	Mean defining features of genome				Mean fitness scores		
Type	Number	Length	Height	Spacing	Overall	Directivity	VSWR
A	5	0.4635	1.7094	0.6957	310.03	316.35	1.32
B	6	0.5047	1.7094	0.5689	308.35	316.89	2.54
C	7	0.5222	1.7130	0.4850	307.39	317.39	3.00
D	7	1.2827	1.7394	0.4740	307.29	346.89	16.30
E	8	1.2588	1.7342	0.4269	307.37	350.66	17.65
F	6	1.3267	0.7589	0.5384	306.72	336.12	11.70
G	9.33	1.2636	1.7244	0.3710	306.15	353.42	18.97

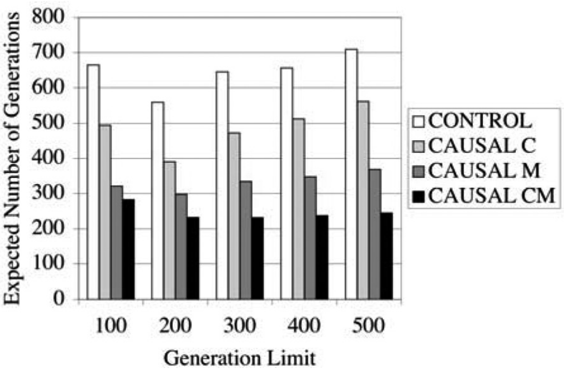

Fig. 8. Expected number of generations required by each system to find an optimal antenna design when used in conjunction with a multi-start strategy, with varying generation limits.

quency with which the different systems converged to antenna designs in the seven categories. The distribution of categories is illustrated in Fig. 9. Of the 200 CONTROL system trials, 123 converged on an optimal category-A antenna design, compared to 139, 177, and 199 of the $\mathrm{CAUSAL_C}$, $\mathrm{CAUSAL_M}$, and $\mathrm{CAUSAL_{CM}}$ trials, respectively. The causally-guided system trials converged to category D, E, and F with less frequency than the control system trials. All 200 of the $\mathrm{CAUSAL_{CM}}$ trials avoided D, E, and F and all but one (category-C) converged to an optimal category-A solution. It is worth noting that the categories that causal systems avoided (D, E, and F) typically have dipole lengths that are longer than optimal and VSWR values that are so high as to be unusable. This relates directly to the user-supplied causal knowledge employed by the causal evolutionary systems in these simulations supporting the hypothesis that the causal relations are contributing effectively to the design process.

To explore the issue of computational cost per generation, 5 trials of the CONTROL and $\mathrm{CAUSAL_{CM}}$ systems were executed for 5000 generations and the CPU time required for each trial was measured. Surprisingly, it was found that despite the expected increased costs associated with the causally-guided systems, on average the CONTROL system required more than 7 times as much CPU time as the $\mathrm{CAUSAL_{CM}}$ system (5893 seconds and 793 seconds, respectively). This difference is largely due to differences in the types of antenna designs that are explored by the two systems. Specifically, the CONTROL system tends to explore antenna designs with longer and more numerous dipole lengths than the causally-guided systems. These types of antennas are more computationally expensive to simulate during fitness evaluation than smaller antennas. Thus, the increased computation per generation required by the $\mathrm{CAUSAL_{CM}}$ system to provide causal guidance is dwarfed by the computational savings of evaluating smaller antennas.

An attempt was made to incorporate an alternative causal relation into the evolutionary system, but this failed to improve performance. Specifically, antenna design experts indicate that there is a strong cause-effect relationship between the height of dipoles and the directivity of an antenna. However, causal relations to that effect did not improve the performance of the system. In an attempt to better understand the cause-effect relationships in this domain the single most highly fit antenna of all trials described above was identified and experimented upon. Four separate experiments were conducted. In each experiment, three of the four antenna genes were held fixed, while the fourth was incrementally adjusted, and changes in antenna performance were observed, yielding insight into causality in this domain and presumably more generally.

Figure 10 illustrates the effects that changing the optimal antenna's number of dipoles, length of dipoles, height of dipoles, and spacing between dipoles has on that antenna's VSWR and directivity. In each graph, a thick black hash on the horizontal axis indicates the unmodified value of the optimal antenna. Consistent with the causal knowledge employed by our system, it can be seen that the length of dipoles has a large effect on the VSWR of an antenna. Noting that the right vertical axis labels of the top right plot in Fig. 10 differ from those of the other plots shown here, the range of VSWR

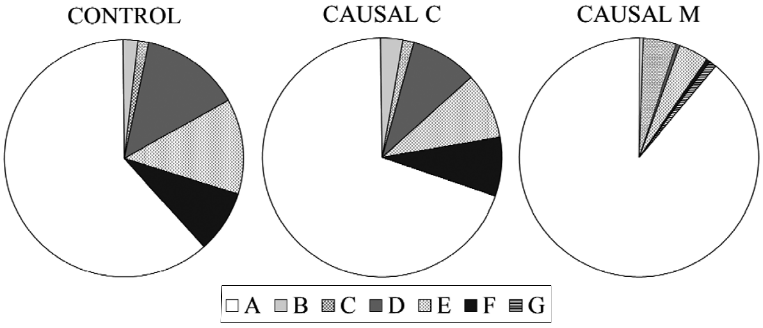

Fig. 9. Distribution of the categories of antenna designs to which each system's trials converge. All but one $\text{CAUSAL}_{\text{CM}}$ trials (not shown) converged to category-A designs.

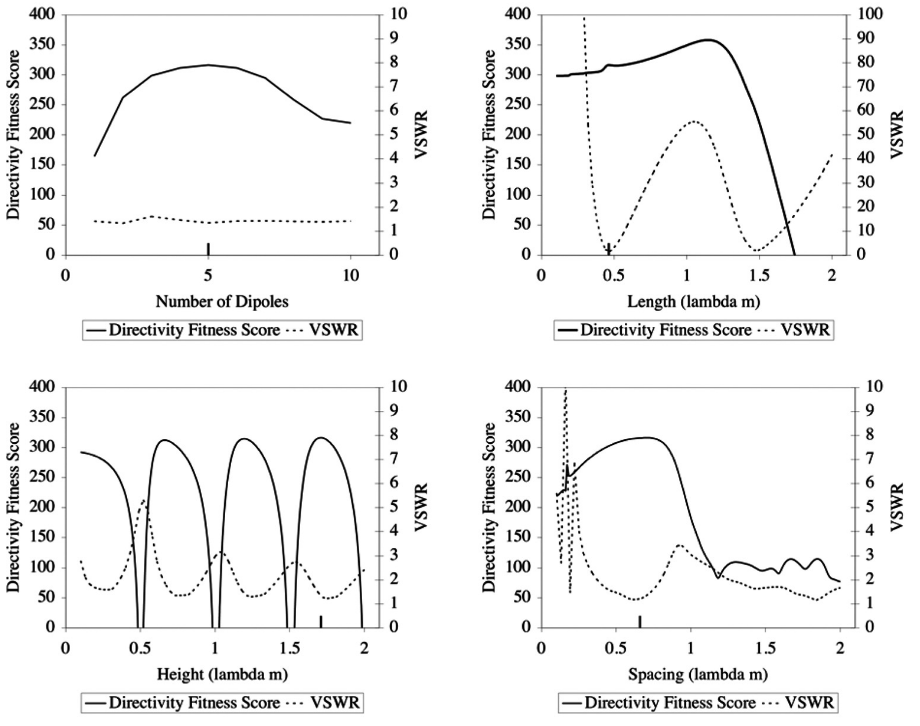

Fig. 10. Effect that changing the number (top left), length (top right), height (bottom left), and spacing (bottom right) has on Fitness Directivity score (left axis) and VSWR (right axis) of the optimal antenna. The range of VSWR values plotted in the top right is an order of magnitude larger than the other three plots.

values when dipole lengths are varied are seen to be an order of magnitude larger than when number, spacing, or height of dipoles are varied. Figure 10 also reveals a clear and seemingly cyclical causal relationship between the height of dipoles and the directivity of antennas. It is clear that the length of dipoles and the spacing between dipoles also has an effect on the directivity of an antenna, but it is difficult to characterize these interactions. These results indicate the complexity of causal relationships in this domain.

4. Discussion

The use of evolutionary computation methods as a design tool to support human engineers has been in part encouraged by the incredible innovativeness of biological evolution processes (e.g., the "invention" of optical lenses, sonar, pumps, valves, winged flight, neural computation, and many other things long before they were thought of by people [9].) Recently there has been growing use of evolutionary computation in engineering, such as the design of control mechanisms for robots [35], mechanical systems [22,30], desalination systems [18], electrostatic micro-generators [20], as well as those discussed in Section 1. In this current work we take some initial steps in evaluating the effectiveness of evolutionary computation methods that have been modified so that explicit cause-effect relations are used to guide the application of genetic operators. Our main innovation is based on the recognition that conventional genetic operators are blind and random, i.e., they execute without regard to the performance characteristics of the individuals to which they are applied.

While our hypothesis was that introducing causally-guided genetic operators would ultimately make evolutionary systems more effective and efficient, there was no guarantee of this a priori; it was entirely possible that just the opposite would be true. Instead, we found that causally-guided evolutionary systems could be used successfully to design dipole antenna arrays that meet pre-specified performance criteria. The performance of these systems was compared to carefully matched control systems that do not employ causally-guided genetic operators. It was found that, at various generations, the causal systems found the fittest antennas with significantly greater frequency than the control system. On average, the causally-guided systems also required significantly fewer generations to find antenna designs with various fitness scores. The causally-guided systems found optimal antenna designs much more frequently largely by avoiding specific sub-optimal designs. Interestingly, these sub-optimal designs were characterized by dipoles that are longer than optimal and have high VSWR values, factors that relate directly to the specific cause-effect relations that the causally-guided system employed. In each result discussed, it was found that the systems using only causal mutation or causal crossover outperformed the control system, but that the system employing both causal mutation and causal crossover performed even better, indicating that these causally-guided operators were synergistic/complementary.

It was found that, surprisingly, the causally-guided system required 1/7th as much CPU time per generation as the control system. This unexpected result was due to differences in the characteristics of antennas that were explored by the various systems, and the different computational costs of simulating those antennas. While this a promising result for causally-guided evolutionary computation in this particular domain, this result is of limited interest as it is domain-specific and may not be relevant to causally-guided evolutionary computation in general. However, it does demonstrate the possibility that in some domains the computational costs of causal inference will be dwarfed by the overall computational costs of the evolutionary systems.

The fact that causally-guided systems were able to solve the dipole antenna array design problem with greater frequency than the control systems demonstrates that using expert-supplied cause-effect relations to bias genetic operations can have a meaningful positive impact on the evolutionary process. The tremendous computational savings of using causally-guided systems with a multi-start strategy are particularly convincing, as this is a very practical measure of computational cost. The fact that the causally-guided system avoids local optima that are directly related to the supplied causal relations is especially encouraging, as this suggests that the causal knowledge is successfully steering the search process away from local optima, much as a human designer might do. It is also encouraging that incorporating either causally-guided mutation or causally-guided crossover into the evolutionary process results in improved performance, and that incorporating both results in even greater performance improvements, suggesting that the value of causal guidance is not critically dependent upon the specific genetic operator being used.

The additional experiments in which design aspects of the fittest antenna were systematically varied and changes in performance measured validated our belief in a causal relationship between dipole length and antenna VSWR. Furthermore, the influence of dipole length on VSWR was found to be an order of magnitude stronger than the influence of any other design aspect on VSWR (at least in the vicinity of an optimum). It is clear from these experiments that there are other causal relationships in the domain, but that such causality can be quite complex. (e.g., a cyclical relationship between dipole height and antenna directivity). This arises because some of the energy that radiates from an antenna is reflected off of the ground plane and passes back over the antenna. There are certain dipole heights at which these reflected waves react destructively with energy radiating directly from the antenna. Unlike the effect of dipole length on VSWR, it appears that the effects of design aspects on directivity are of similar magnitude. These results help to explain why preliminary efforts to include a causal relationship between dipole height and antenna directivity into the evolutionary system were unsuccessful, while the causal relationship between dipole length and VSWR was successful.

They suggest that with our current approach very strong and straightforward relationships may be incorporated, while more complex relationships involving multiple design aspects may be problematic.

The causally-guided evolutionary system presented here can be viewed in a number of ways. First, is as a novel form of adaptive parameter control. In adaptive parameter control methods, as the evolutionary process runs, various characteristics of the process are monitored, statistics are computed, and parameters are adjusted according to pre-defined heuristics [15]. However, we know of no previous adaptive parameter control method in which the performance characteristics of evolved individuals are examined using cause-effect relations that bias the execution of the genetic operators. Additionally, while the vast majority of adaptive parameter control methods make dynamic adjustments to parameters that are global to the entire population, our causally-guided evolutionary system falls under a much smaller category of adaptive methods in which parameters are adjusted per individual [23]. Second, our approach relates to past studies in fitness approximation [24], where an approximate model of the fitness function is used to efficiently estimate the fitness of individuals. In some past work, the fitness approximations of individuals have been used to guide genetic operators [1,39], in a similar fashion to the work presented here. However, in causally-guided evolution there is no fitness approximation (the actual fitness is calculated) and instead the causal knowledge is used to guide the genetic operators. Third, causally-guided evolution can be viewed as being an instance of a memetic algorithm. The term memetic algorithm is sometimes used to refer to a wide range of approaches in which multiple optimization strategies are integrated in to a single approach [37], such as combining evolutionary computation with some form of local search [13,26,28, 36]. In the case of causally-guided evolutionary computation, conventional evolutionary search is combined with causal inference. To our knowledge, no previous studies involving memetic algorithms have used explicit causal relations as the basis for local search.

Finally, there remain some important areas for future study of causally-guided evolutionary computation. In order to further evaluate the methods introduced here, causally-guided evolutionary computation should be applied to additional challenging applications. While a preliminary study has previously been successfully applied to a "toy" neural network design problem [11], more complex and varied application domains should be explored and further analyzed statistically. One particularly important issue for future applications is what impact causally-guided genetic operators have on discovering truly novel solutions. Another is to assess the practical significance of the improved computational efficiency that we observed, such as with evolution of adaptive designs in which continual re-evolution is required. Still another is to compare the results we obtained here against those obtainable with other optimization methods (PSO, simulated annealing, etc.). Further, an important extension of this work would be to expand the ways in which human experts may describe causal knowledge, including non-linear relationships and their associated probabilities.

5. Conclusions

We conclude that causally-guided genetic operators offer an effective way of incorporating expert-supplied cause-effect domain knowledge into the evolutionary design of antennas. Causally-guided evolution appears to increase the efficiency and effectiveness of the evolutionary process, by directing the evolutionary search away from less fruitful areas of the search space.

Acknowledgment

This work was supported by NSF Award IIS 075 3845.

References

[1] K. Abboud and M. Schoenauer, Surrogate deterministic mutation, in: *Artificial Evolution '01*, Springer, 2002, pp. 103–115.
[2] H. Adeli and N.-T. Cheng, Augmented Lagrangian genetic algorithm for structural optimization, *Journal of Aerospace Engineering* **7**(1) (1994), 104–118.
[3] H. Adeli and N.-T. Cheng, Concurrent genetic algorithms for optimization of large structures, *Journal of Aerospace Engineering* **7**(3) (1994), 276–296.
[4] H. Adeli and N.T. Cheng, Integrated genetic algorithm for optimization of space structures, *Journal of Aerospace Engineering* **6**(4) (1993), 315–328.
[5] E. Altshuler and D. Linden, Wire-antenna designs using genetic algorithms, *Antennas and Propagation Magazine IEEE* **39**(2) (1997), 33–43.
[6] T. Back, U. Hammel and H. Schwefel, Evolutionary computation: Comments on the history and current state, *IEEE Trans Evolutionary Computation* **1** (1997), 3–17.
[7] P. Baraldi, C. Roberto, Z. Enrico, S. Redouane and R. Chevalier, Genetic algorithm-based wrapper approach for grouping condition monitoring signals of nuclear power plant components, *Integrated Computer-Aided Engineering* **18**(3) (2011), 221–234.

[8] R. Becerra and C.C. Coello, A cultural algorithm for job-shop scheduling, in: *Knowledge Incorporation in Evolutionary Computation*, Y. Jin, ed., Springer, 2005, pp. 37–56.

[9] P.J. Bentley, An introduction to evolutionary design by computers, in: *Evolutionary Design by Computers*, P.J. Bentley, ed., Morgan Kaufmann, 1999, pp. 1–73.

[10] G.J. Burke and A.J. Poggio, *Numerical electromagnetics code NEC: Method of moments*, Lawrence Livermore Lab, Tech. Rep. UCID18834, January 1981.

[11] T. Chabuk and J. Reggia, Causally-guided evolutionary computation for adapting weights, in: *Proc 4th Intl Conf on Neural, Parallel and Scientific Computations*, Atlanta GA, Dynamic Publishers, 2010, pp. 77–82.

[12] C.M. Coleman, E.J. Rothwell and J.E. Ross, Investigations of simulated annealing, ant-colony optimization, and genetic algorithms for self-structuring antennas, *IEEE Trans Antennas and Propagation* **52**(4) (2004), 1007–1014.

[13] K. DeJong, *Evolutionary Computation*, MIT Press, 2006.

[14] S. Drabowitch and A. Papiemik et al., *Modern Antennas*, Chapman and Hall, 1998.

[15] A. Eiben, Z. Michaelwicz, M. Schoenauer and J.E. Smith, Parameter control in evolutionary algorithms, in: *Parameter Setting in Evolutionary Algorithms*, F. Lobo, C. Lima and Z. Michaelwicz, eds, Springer, 2007, pp. 19–46.

[16] R. Elliot, *Antenna Theory and Design*, IEEE Press, 2003.

[17] D. Fogel, What is evolutionary computation? *IEEE Spectrum* **37** (2000), 26–32.

[18] K. Hamza, M. Shalaby, A.O. Nassef, M.F. Aly and K. Saitou, Design optimization of reverse osmosis water desalination systems via genetic algorithms, in: *ASME Conf. Proceedings*, 2010, pp. 655–664.

[19] R.L. Haupt and D. Werner, *Genetic Algorithms in Electromagnetics*, IEEE Press, 2007.

[20] D. Hoffmann, B. Folkmer and Y. Manoli, Multi-parameter optimization of electrostatic microgenerators using design optimization algorithms, *Smart Materials and Structures* **19**(11) (2010), 115016.

[21] G. Hornby, A. Globus, D. Linden and J. Lohn, Automated antenna design with evolutionary algorithms, in: *Proc American Institute of Aeronautics and Astronautics Space*, 2006, pp. 1–8.

[22] J. Hu, E.D. Goodman, S. Li and R. Rosenberg, Automated synthesis of mechanical vibration absorbers using genetic programming, *Artificial Intelligence Engineering Design Analysis Manufacturing* **22**(3) (2008), 207–217.

[23] H. Iba and H. de Garis, Extended genetic programming with recombinative guidance, in: *Advances in Genetic Programming 2*, P.J. Angeline and K.E. Kinnear, eds, Cambridge MA, MIT Press, 1996.

[24] Y. Jin, A Comprehensive Survey of Fitness Approximation in Evolutionary Computation, *Soft Computing* **9**(1) (2005), 3–12.

[25] Y. Jin, ed., *Knowledge Incorporation in Evolutionary Computation*, 2005, Spring Verlag.

[26] J. Knowles and D. Corne, Memetic algorithms for multiobjective optimization: Issues, methods and prospects, in: *Recent Advances in Memetic Algorithms, ser. Studies in Fuzziness and Soft Computing*, (vol. 166), W. Hart, J. Smith and N. Krasnogor, eds, Springer Berlin/Heidelberg, 2005, pp. 313–352.

[27] J.R. Koza, *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*, Kluwer Academic Publishers, 2003.

[28] N. Krasnogor and J. Smith, A tutorial for competent memetic algorithms: model, taxonomy, and design issues, *IEEE Trans Evolutionary Computation* **9**(5) (2005), 474–488.

[29] A. Kusiak and H. Zheng, Optimization of wind turbine energy and power factor with an evolutionary computation algorithm, *Energy* **35**(3) (2010), 1324–1332.

[30] H. Lipson, Evolutionary synthesis of kinematic mechanisms, *Artificial Intelligence Engineering Design Analysis Manufacturing* **22**(3) (2008), 195–205.

[31] J. Lohn, G. Hornby and D. Linden, Human-competitive evolved antennas, *Artificial Intelligence Engineering Design Analysis Manufacturing* **22**(3) (2008), 235–247.

[32] J. Lohn, W. Kraus and D. Linden, Evolutionary optimization of a quadrifilar helical antenna, in: *Antennas and Propagation Society Intl. Symposium*, IEEE, 2002, pp. 814–817.

[33] J. Lohn, W. Kraus, D. Linden and S. Colombano, Evolutionary optimization of yagi-uda antennas, in: *Proc. 4th Intl. Conf. on Evolvable Systems: From Biology to Hardware* (*ICES '01*), London, Springer-Verlag, 2001, pp. 236–243.

[34] J. Lohn, D. Linden, G. Hornby and W. Kraus, Evolutionary design of an x-band antenna for nasa's space technology 5 mission, in: *Antennas and Propagation Society Intl. Symposium*, IEEE, 2004, pp. 2313–2316.

[35] O. Michel, Evolutionary neurogenesis applied to mobile robotics, in: *Advances in the Evolutionary Synthesis of Intelligent Agents*, M. Patel, V. Honavar and K. Balakrishnan, eds, MIT Press, 2001, pp. 185–213.

[36] P. Moscato and C. Cotta, A modern introduction to memetic algorithms, in: *Handbook of Metaheuristics, ser. Intl. Series in Operations Research and Management Science*, M. Gendreau and J.-Y. Potvin, eds, Springer, 2010, pp. 141–183.

[37] F. Neri, C. Cotta and P. Moscato, *Handbook of Memetic Algorithms*, Springer, 2011.

[38] E. Rajo-Iglesias and O. Quevedo-Teruel, Linear array synthesis using an ant colony optimization based algorithm, *Antennas and Propagation Magazine* **49**(2) (2007), 70–79.

[39] K. Rasheed, X. Ni and S. Vattam, Methods for using surrogate models to speed up genetic algorithm optimization: informed operators and genetic engineering, in: *Knowledge Incorporation in Evolutionary Computation*, Y. Jin, ed., Springer, 2005, pp. 37–56.

[40] J. Robinson and Y. Rahmat-Samii, Particle swarm optimization in electromagnetics, *IEEE Trans Antennas and Propagation* **52**(2) (2004), 397–407.

[41] A. Rodriguez and J. Reggia, A Distributed Learning Algorithm for Particle Systems, *Integrated Computer-Aided Engineering* **16**(1) (2009), 1–20.

[42] L. Setiean, *Practical Communication Antennas*, Prentice Hall, 1998.

[43] J. Smith and T. Nguyen, Autonomous and cooperative robotic behavior based on fuzzy logic and genetic programming, *Integrated Computer-Aided Engineering* **14**(2) (2007), 141–159.

[44] D. Whitley, Permutations, in: *Evolutionary Computation 1*, T. Back, D. Fogel and Z. Michalewicz, eds, IOP Press, 2000, pp. 139–150.

[45] J. Zhang, H. Chung, A. Lo and B.J. Hu, Fuzzy Knowledge Incorporation in Crossover and Mutation, in: *Knowledge Incorporation in Evolutionary Computation*, Y. Jin, ed., Springer, 2005, pp. 37–56.

[46] L. Zhang, R. Raut, J. Yingtao, K. Ulrich and K. Yoohwan, Macro-cell placement for analog physical designs using a hybrid genetic algorithm with simulated annealing, *Integrated Computer-Aided Engineering* **12**(4) (2005), 379–396.