The Use of Nonparametric Methods and Evolutionary Algorithms in Genetic Epidemiology of Complex Disease

By Colleen M. Farrelly

1) INTRODUCTION

Technological advances in genome sequencing of populations and families have provided geneticists and epidemiologists with a wealth of resources to aid in the exploration of complex disease etiology. However, these advances are fraught with many analytical challenges that must be addressed if researchers are to make full use of these resources.

Obtaining the power necessary to detect risk factors contributing to an increased disease incidence of only 1.3-fold or less within a large dimensional dataset consisting mainly of noise presents a significant challenge (Moore and Williams, 2009). Commercially-available genotyping, such as the chips designed by Affymetrix and Illumina, can tag over 500,000 single nucleotide polymorphisms (SNPs), and the advent of newer, faster sequencing methods may increase this number in the future (Klein, 2007). Klein's power analysis of these genotyping studies suggest that the minimum number of individuals needed to find a genotypic relative risk of 1.5 at 80% power is around 3,500, depending on the sequencing methods (Klein, 2007). Recruitment and current sequencing costs may limit researchers' abilities to find low risk or rare variants associated with disease in new populations, though several databases include genome sequencing data from an adequate number of individuals. Traditional parametric methods of analysis, such as logistic regression, do not have enough power to detect main effects and interactions in such datasets, which usually violate methodological assumptions about the data; semiparametric or nonparametric techniques, such as random forests (RF), multifactor dimensionality reduction (MDR), and genetic programming optimized neural networks (GPNN), are necessary to provide the power needed to identify risk SNPs (Heidema et al., 2006).

In addition to power challenges, large numbers of independent variables relative to sample size, commonly referred to as "the curse of dimensionality," also restrict the use of certain methods of analysis. Parametric methods of analysis and imputing missing data, such as Markov Chain Monte Carlo multiple imputation, require more participants than independent variables and, thus, cannot be used without reducing the number of predictors prior to analysis (Heidema et al., 2006; Gheys & Smith, 2009). The large volume of data also limits the use of certain nonparametric methods by increasing computing time to unfeasible levels. For instance, combinatorial methods used to detect multi-way gene-gene interactions (epistasis) and gene-environment interactions (plastic reaction norms), such as MDR or restrictive partitioning methods, collapse data into smaller numbers of groups based upon evaluation of all possible n-way variable combinations, resulting in large Bonferroni corrections for multiple tests and computational limits, as the number of interaction terms

searched grows exponentially as the number of possible predictors increases (Culverhouse et al., 2004; Bush et al., 2006). Attribute selection methods, such as the ReliefF filter approach or stochastic search wrapper approaches, ameliorate some of this computational burden, but such methods can lead to problems of underfitting and overfitting models, as well as introducing another source of error into models (Moore et al., 2010; Han et al., 2004).

Further, it is thought that epistasis, gene-gene interactions without strong main effects, and plastic reaction norms, the gene-environment analog of epistasis, play an important role in the development of complex diseases, as many genome-wide association studies (GWAS) searching for main effects have not found SNPs that account for significant portions of variance and are sometimes not replicable by future studies (Culverhouse et al., 2004, Moore & Williams, 2009; Heidema et al., 2006; Moore et al., 2010). Rare variants, low penetrance, interactions likely complicate the analysis of complex diseases, as opposed to the relatively-simple case of Mendelian disease (Moore & Williams, 2009). Biologically, epistasis and plastic reaction norms can be explained through molecular interactions in biochemical pathways and through epigenetic changes in chromosome structure affecting gene expression, respectively (Greene et al., 2009; Lou et al., 2007; Moore and Williams, 2009; Moore et al., 2010). For example, in addiction, genetic and environmental factors (such as repeated exposure to a drug) interact biochemically to change histone structure of transcription factor genes (CREB, Δ FosB, NF-KB, MEF-2, and EGRs) through methylation, phosphorylation, and acetylation, making some genes more likely and others less likely to be transcribed within a cell (Robison & Nestler, 2011). Statistically, both represent nonadditive effects in linear models (Moore & Williams, 2009), which, in the absence of main effects, seriously limits the use of parametric techniques (Heidema et al., 2006). However, many nonparametric techniques thrive in this situation and were, in fact, developed for such a situation (Moore & Williams, 2009; Heidema et al., 2006).

Along with interactions within a biological pathway, complex diseases often involve multiple pathways, a phenomenon known as genetic heterogeneity. For example, opiate addiction has been shown to involve the brain's dopaminergic, noradrenergic, and endogenous opioid pathways (Robison & Nestler, 2011). Methods robust to many of the challenges posed by genomic data, such as combinatorial methods and set association, often aim to find an optimal solution, rather than several significant solutions, thereby missing important contributions to variance (Heidema et al., 2006; Pattin et al., 2009).

Related to genetic heterogeneity is the phenomenon of phenocopies, individuals with low genetic risk, who, nevertheless, develop the disease of interest. Phenocopies decrease the association of risk genes in different pathways involved in the disease process with the development of disease and pose significant problems in genetic epidemiology (Heidema et al., 2009). Including environmental factors as independent variables can reduce the impact of phenocopies on identifying risk SNPs and provide a more comprehensive picture of disease etiology.

The last statistical challenge facing genetic epidemiology is multicollinearity. Genes physically close to each other on chromosomes show different inheritance patterns than genes further apart, known as linkage disequilibrium (Ziegler et al., 2008). Using haplotypes, clusters of genes in linkage disequilibrium with each other and not likely to crossover during meiosis, rather than SNPs, in analyses, as well as adjusting gene importance measures, has shown promise in alleviating the bias in results (Ziegler et al., 2008; Meng et al., 2009).

2) ANALYTIC TECHNIQUES

2.1) Parametric and Semiparametric Techniques

Logistic regression is a common type of regression in which predictors, such as SNPs and environmental factors, are linked to a binary outcome variable via a logit function. Significant variables are added to a model with forward selection, which can also involve interaction terms provided a main effect exists, or a full model can be pruned with backward selection (Heidema et al., 2006). Another procedure, called least absolute shrinkage and selection operator (LASSO), may be employed to shrink coefficients of unimportant variables to 0, thereby reducing model size; however, this method, like forward and backward selection, suffers when a large number of predictors relative to sample size are present or in the presence of multicollinearity or genetic heterogeneity (Heidema et al., 2006). The employment of evolutionary algorithms, such as genetic algorithms, has proven to be an effective method of variable selection in multiple regression, as well as logistic regression, and may represent a potential solution to some of the problems arising in this technique with respect to genetic epidemiology (Najafi et al., 2011; Gayou et al, 2008; Broadhurst et al, 1997; Paterlini & Minerva, 2010).

Artificial neural networks (NNs) represent a hybrid of parametric and nonparametric techniques. NNs utilize a directed graph of connected node layers in an optimum architecture to process data and detect underlying patterns (Motsinger-Reif et al., 2008). In traditional multilayer perceptrons, an input node layer receives predictors in a data set, which is then processed by one or more hidden node layers of "transfer functions," such as logistic regression, before exiting through the output layer, which is used to classify the information into the dependent variable's categories or range (Motsinger-Reif et al., 2008; Heidema et al., 2006). Each connection between nodes is assigned an adjusted weight of its transfer function through backpropagation as the NN is trained on a cross validation bootstrap sample of a data set; error estimates are then obtained thorough a test set (Heidema et al., 2006; Venayagamorthy & Singhal, 2005). Increasing the number of hidden layers and nodes in those layers allows a NN to capture complex, nonlinear relationships and interaction effects among input variables (Heidema et al., 2006). Similar to classical multilayer perceptrons, simultaneous recurrent NNs employ a context feedback layer within their input layer, which receives the NN output,

to aid in computationally-complex processing (Venayagomorthy & Singhal, 2005).

However, in complex training data, such as those encountered in genetic epidemiology, the backpropagation algorithm can stall in local minima, leading to suboptimal fit and performance; exhaustive search through all possible configurations of a NN architecture is computationally prohibitive (indeed, sometimes impossible), as even small NN's potential solution would have a run time of many years (Motsinger-Reif et al., 2008). NNs are also limited in the number of input variables they can process, creating variable selection problems in large data sets, such as genomics data (Heidema et al., 2006).

Evolutionary computing/algorithms, such as genetic programming and grammatical evolution (both of which use a genetic algorithm to evolve computer programs to an ideal program to solve a particular problem, such as NN structure optimization), has shown promise in drastically reducing computing time while arriving at globally-optimum solutions (Ritchie et al., 2003; Motsinger-Reif et al., 2008; Zhou et al., 2001). These methods have yielded promising results in the analysis of simulated and real genomics data sets, and grammatical evolution, in particular, has proven computationally tenable for use in datasets containing >500,000 SNPs (Motsinger-Reif et al., 2008). Another technique involves NN ensembles evolved through a genetic algorithm (Zhou et al., 2001), which shows similar performance on UCI repository datasets to other ensemble methods (such as random forests). A promising new development, which has yet to be tested on real-world datasets, is the use of quantum evolutionary algorithms (refer to section 3.2) in place of backpropagation to train multilayer perceptrons and simultaneous recurrent NNs, which are computationally challenging and expensive to train (Venayagamoorthy & Singhal, 2005). Mean square errors were better than traditional training methods, especially with simulated complex, noisy data, and computational times were dramatically reduced (Venayagamoorthy & Singhal, 2005). However, this study employed pre-specified NN structures, which are unknown a priori in most real-world situations, and did not test this method with datasets similar to those encountered in genetic epidemiology.

2.2) Nonparametric Methods

2.2.1) Cluster and Combinatorial Classification Methods

2.2.1.1 Cluster Methods

Two group distance-based approaches are the K-means clustering algorithm and the K-nearest neighbors (KNN) approach. The K-means clustering (KMC) algorithm iteratively partitions its dataset's N-dimensional space, optimizing outcome similarities of data points assigned to the same hyperplane partition and outcome differences of points in different partitions through distance metrics, i.e. minimizing within-cluster distance while maximizing between-cluster distance (Xiao et al., 2008; Maulik & Bandyapadhyay, 2000). Generally, this method deals well with massive datasets. Pairing KMC with evolutionary algorithms, such as a quantum-inspired genetic algorithm,

improves speed and accuracy in small and medium-sized datasets; however, these have yet to be tested on datasets on the scale of genomics data (Xiao et al., 2008).

The similar KNN has been used extensively in the classification of microarray data, which suffers from some of the same problems facing genetic epidemiology (Li et al., 2001). This approach considers each data point in the context of its k nearest neighbor points, as measured by geometric distance in space, such as Euclidean distance or geometric mean distance (Li et al., 2001; Lee et al., 2005). If the k -nearest neighbors have the same classification group, a point is classified into that group; if not, a point is considered to be unclassifiable (Li et al., 2001; Jirapech-Umpai & Aitken, 2005). While KNN accommodates interactions and genetic heterogeneity, massive datasets, including larger microarrays substantially smaller than genome-wide datasets, present computational challenges (Li et al., 2001; Ooi & Tan, 2003; Heidema et al., 2006). Several attempts have been made to reduce the number of parameters and to optimize variable selection for KNN approaches through the use of genetic algorithms; testing results on the Golub et al. Leukemia dataset, containing 7129 genes from 72 individuals, yield correct prediction rates of 92% (Deutsch, 2003, GESSES algorithm), 97% (Jirapech-Umpai & Aitken, 2005, RankGene algorithm), and 61% (Li et al., 2001, GAKNN). Opportunities exist in the development of KNN with more powerful, computationally feasible evolutionary algorithms, such as quantum-inspired evolutionary algorithms.

2.2.1.2) Combinatorial Methods

Combinatorial methods, which include combinatorial partitioning (CPM), restrictive partitioning (RPM), and MDR, identify combinations of variables explaining large chunks of variance (epistasis and plastic reactive norms) by searching through all possible combinations of predictor variables, which may include SNPs or environmental factors (Heidema et al., 2006), and evaluating their ability to predict outcomes. CPM performs an exhaustive search for the best n -way interactions out of a given collection of p variables, searching through C_n^p possible solutions and validating selected sets through multifold cross validation (Heidema et al., 2006; Culverhouse et al., 2004). For large datasets, computational limits and CPM's multiple testing design necessitate directed search techniques or variable selection to reduce dimension before analysis.

To deal with multiple testing problems and computational challenges posed by CPM, Culverhouse et al. (2004) developed RPM, which selectively searches through possible purely epistatic models to find the optimum combination as determined by a model's R^2 value. This algorithm iteratively merges similar genotypes and partitions data into good combination areas for further exploration and bad areas to be avoided in future searches (Culverhouse et al., 2004). In simulation studies, the method has proven accurate, and RPM has been successfully employed in real datasets, as well (Culverhouse et al., 2004). However, this method still cannot handle large datasets

computationally, and it suffers from multiple testing issues, which limit RPM's ability to detect significant effects (Heidema et al., 2006).

MDR is a more widely-used combinatorial method, which has been successfully developed for both population-based studies and pedigree-based studies (Bush et al., 2008), and has been proven to be the best method for indentifying multilocus epistasis (Hahn et al., 2002). In this method, data are divided into a training set and a test set, and the training set is then evaluated for possible n-way combinations. A case-control ratio threshold, usually set to 1, is chosen, and combinations are assigned to high-risk (>1) or low-risk (<1) categories based on a particular genotype's case-control ratio (kernels G_1 and G_0 , respectively). Combination errors are then calculated for each n-order pair, and the best n-order combination is chosen for prediction error evaluation and cross validation by testing set data. The best of the n-order models is selected for permutation testing to confirm the contribution of each of the n genes in the model (Bush et al., 2008; Lee et al., 2007; Lou et al., 2007).

While this method shows promise, it suffers from several problems, including inability to process large datasets, difficulties related to missing combinations of genotypes in a given dataset, and problems when faced with genetic heterogeneity, as only one combinatorial model of interactions is identified by MDR (Greene et al, 2009; Lou et al., 2007; Lee et al., 2007). To handle large data sets more effectively, two techniques have been developed recently: parallel MDR (Bush et al., 2006) and variable selection methods (Moore & Williams, 2009). Parallel MDR relies on a tree-based recursive binning technique, allowing for more efficient data handling, model generation and processing, and storing of solutions; this method has proven effective for datasets of hundreds of thousands of SNPs with high-order ($n>5$) interaction terms (Bush et al., 2006). However, different strategies of model evaluation are likely necessary to deal with genetic heterogeneity and computational cost of permutation testing.

A more commonly used approach to computational challenges is variable selection prior to MDR application. This can be accomplished through the use of filter methods, which rely on machine learning strategies, or of wrapper methods, which utilize probabilistic stochastic search algorithms (Moore & Williams, 2009; Greene et al., 2009). Historically employed filters have included variations of the Relief algorithm, which examines a data point's nearest neighbor, one with the same outcome (a hit) and one with a different outcome (a miss), and scores that individual as a potential outcome predictor. Variations include ReliefF, which considers multiple nearby hits and misses; Tuned ReliefF (TuRF), which iteratively deletes SNPs with low ReliefF scores; Spatially-Uniform ReliefF (SURF), which searches all hits and misses within a finite radius of an individual point; and SURF & TuRF, which combines the SURF algorithm with the iterative deletion method of TuRF (Greene et al., 2009). Of these methods, SURF & TuRF has been shown to be the most effective and efficient filter approach to MDR, handling large data sets, low heritability, and small effect sizes (Greene et al., 2009). Though wrapper approaches, such as

genetic programming, simulated evaporative cooling, and particle swarm optimization, offer another effective approach, little work has been done in this area to date (McKinney et al., 2009).

To deal with empty genotype combinations and possible interacting covariates, two methods have been developed to assuage computational issues of these problems. First, Lee et al. (2007) have proposed and tested a log-linear model-based MDR, in which a saturated model (at least 1 individual matching each possible combination of n chosen variables) corresponds to the familiar MDR method. This method provides more power and smaller error rates when confronted with empty genotype combinations in data sets than the usual MDR method (Lee et al., 2007). A generalized MDR, based upon generalized linear models consisting of linking functions (such as the identity function or logit function), interacting variables, covariates, and variable-covariate interactions has also been developed as a more flexible, more comprehensive approach to MDR; it has proven effective at dealing with noise and differently-scaled variables in a real-world nicotine dependence dataset (Lou et al., 2007).

An interesting new development offering the possibility of indentifying all significant n -way interactions by MDR employs hypothesis testing via an extreme value distribution (EVD), rather than expensive permutation testing (Pattin et al., 2009). This method is 50 times faster than 1000-fold permutation testing and is robust to heritability and sample size variations without sacrificing performance accuracy. No differences in chosen EVDs were noted, suggesting a possible extension with EVDs better equipped to handle linkage disequilibrium and main effects, which violate assumptions of the generalized EVD used in this study (Pattin et al., 2009).

2.2.2) Tree-Based Methods

Tree-based classification methods have proven useful in the analysis of microarray and GWAS data, tackling problems of dimensionality, genetic heterogeneity, and epistasis while maintaining power and accuracy (Heidema et al., 2006; Fou & Gray, 2005; Lunetta et al., 2004; Diaz-Uriarte & Andres, 2006; Bureau et al., 2004).

2.2.2.1) Single-Tree Methods

The simplest and easiest to interpret, albeit less accurate, of regression (continuous data) and classification (categorical data) tree methods are single-tree methods, including classification and regression trees, CART (Breiman et al., 1984); Bayesian CART (Denison et al., 1998; Chipman et al., 1998), and Tree Analysis with Randomly Evolved Trees, of TARGET (Fou & Gray, 2005). The most straight-forward method, CART, builds binary decision trees using predictor variables to form splitting rules (at each branch "node") with respect to an outcome variable (Breiman, 1984; Loh & Shih, 1997). Models are fully-grown and then pruned by backward selection to the best model size (number of terminal nodes, branching nodes, and depth). While a good method, many new methods outperform CART when tested on real-world data sets, such as Servo or Boston Housing, from the UCI Machine

Learning Repository (Fan & Gray, 2005; Breiman, 2001; Denison et al., 1998). However, ensemble methods (which grow and draw inference from multiple trees, usually grown on randomly-selected subsets of predictor variables), such as random forests, bagging, or Adaboost, usually use CART to grow their collections of trees and have shown good results with this method (Breiman, 2001; Hothorn et al., 2004).

Bayesian CART improves upon the CART algorithm by searching through the possible tree space probability distribution through reversible jump Markov Chain Monte Carlo methods using a hybrid sampler to avoid local traps (Denison et al., 1998). This method essentially identifies "fertile" areas of the multivariate tree space probability distribution, which produce good trees. A similar version developed by Chipman et al. (1998) utilizes this knowledge when constructing trees, rather than searching all possible node split rules, and selects the best tree as its output model, allowing for easy visual interpretation. These methods outperform CART when tested on the UCI Air dataset (Fan & Gray, 2005).

TARGET combines single tree methodology with another stochastic search technique, genetic algorithms, to evolve a population of randomly-generated possible regression trees according to genetic operators (see Section 3) until the algorithm converges to an ideal tree, as assessed by the Bayesian Information Criterion (BIC), which is given in the output (Fan & Gray, 2005; Cha & Tappert). The use of BIC as a measure of fit considers prediction accuracy, as well as model complexity, when evaluating possible ideal tree models, aiding in the interpretation and generalization of results. TARGET outperforms both CART and Bayesian CART on the UCI Air dataset, with an average reduction in residual sum of squares values around 5% (Fan & Gray, 2005). On the UCI Boston Housing dataset, TARGET outperforms CART and multiple regression; yields similar mean square error values as neural networks, Bayesian Additive Regression Trees (BART, a Bayesian ensemble technique), and Adaboost (a tree ensemble method using a boosting algorithm); and is outperformed by adaptive bagging, random forests, and bagging (Fan & Gray, 2005; Breiman, 2001; Chipman et al., 2010). On Breiman's Relative Assessment of Tree Modeling Methods, TARGET receives an A- in predictive capability (compared to A+ with RFs, B with CART) and an A++ in interpretability (F with RFs, A+ with CART). This represents a potential new tree-growth mechanism for random forests with massive data sets and a starting point for the use of other evolutionary algorithm-based optimization techniques, such as quantum-inspired evolutionary algorithms, within tree-based methodology.

2.2.2.2) Ensemble Methods

Ensemble-based methods, in which many trees are grown with split rule selection based upon randomly drawn variable subsets, include Adaboost, bagging, BART, RFs, and RF extensions (Breiman, 2001; Chipman et al., 2010; Hothorn et al., 2004; Zhang et al., 2003). These methods have been developed to create greater stability amongst chosen predictors, as single-tree methods may have several near-ideal tree structures based on different variables splitting tree nodes

(Breiman, 2001); such as situation may arise from genetic heterogeneity, where each disease pathway may yield a near-ideal tree in vastly different ways (tree size, variables chosen, structure...). Bagging is a technique suited for data sets in which the importance of predictors is not known a priori and examines overall classification among trees, rather than voting or averaging across trees (Breiman, 2001; Hothorn et al., 2004). However, this method is outperformed by other methods, such as random forests (Breiman, 2001), and does not provide intuitive or interpretable output with respect to selected predictors' contribution to the outcome of interest, an important function of modeling genetic data.

Bayesian Additive Regression Trees, known as BART, is a robust, additive sum-of-trees model of random components with adaptive dimension fit through a Markov Chain Monte Carlo method employing a Metropolis-Hastings algorithm to grow trees based on a prior distribution (Chipman et al., 2010). It is based on a boosting algorithm similar to Adaboost, which utilizes sequences of trees in a similar fashion to the way multiple regression uses sequences of predictor variables, rather than a data randomization and search algorithm, upon which bagging and random forests are based (Chipman et al., 2006; Breiman, 2001). BART's performance on various UCI datasets outperforms single-tree methods, other boosting techniques, and random forests, while handling complex data more quickly and efficiently than other methods (Chipman et al., 2010). Further testing is needed to determine if BART can handle datasets as large as those used in genetic epidemiology, but BART offers a more effective technique that is computationally faster than random forests, which have limited analytical capabilities in very large data sets (Zhang et al., 2009).

RFs, created by Breiman (2001), are ensemble methods utilizing random split selection on split training data, in which different subsets of variables are randomly drawn with or without replacement to determine node splitting rules at a particular node in a maximally-growing tree, and tree voting methods, in which each tree with a given variable contributes to an overall variable importance measure (traditionally Gini Importance with voting or Permutation Importance with permutation testing on out-of-bag observations—i.e. individuals not chosen when building a particular tree). RFs are stable predictors capable of handling interaction effects (connected nodes in a pathway leading to a terminal node, or leaf, containing classification information for that pathway) and large amounts of data (Heidema et al., 2006; Lunetta et al., 2004; Zhang et al., 2009; Meng et al., 2009). RFs converge to solutions absolutely (Breiman, 2001) and do not suffer from overfitting, though fit quality may be poorer with highly correlated predictors (Segal, 2003).

However, RFs' importance measures have suffered from multicollinearity (owing to linkage disequilibrium), as well as bias towards variables with more categories and indirect measures of interactions among variables within trees (Meng et al., 2009; Lunetta et al., 2004). To deal with bias, Strobl et al. (2007) suggest using permutation testing, rather than Gini measures of node purity, which attenuate the bias. To address the issue of linkage disequilibrium and correlated predictor problem in general, Meng et al. (2009)

demonstrate the efficacy of a revised importance measure (rIM) based on selection of splitting variables in linkage equilibrium, which can be employed without also correcting tree-building methods for linkage disequilibrium. Bureau et al. (2004) also address importance measures and suggest a joint-effects framework, which aids in interaction detection and ameliorates bias from multicollinearity amongst predictor variables.

RFs and their derivatives have been used extensively in the classification of microarray data, as well as the analysis of GWAS, including simulations (Diaz-Uriarte & Andres, 2006), asthma (Bureau et al., 2004), age-related macular degeneration (Jiang et al., 2009; Chen et al., 2007), alcoholism (Ye et al., 2005), smoking (Ye et al., 2005), adverse small pox vaccine reactions (McKinney et al., 2009), and various cancers (Dressman et al., 2007; Pittman et al., 2004). Three extension of RFs have been developed recently that show promise in the analysis of genomics data, as well. Enriched RFS, created by Amaratunga et al. (2008), aims to reduce error in data sets with large amounts of noise by weighting known predictors more highly than potential noise variables during subset selections employed in splitting rule evaluation, so as to stack the chances of finding a tree predictor within a randomly-drawn subset. However, this method requires knowledge of potential predictor SNPs and biochemical pathways and, thus, would not be an effective method for identifying previously-unknown risk SNPs.

Deterministic forests (DFs), which grow trees based upon the best n root node splits for tree construction to a predetermined depth, address heterogeneity, reduce prediction error of a forest, and increase external validity of findings (Zhang et al., 2003; Zhang et al., 2009; Chen et al., 2007; Ye et al., 2005). DFs have effectively handled previously published genomics datasets, including the Leukemia and Lymphoma datasets, well and identified new variants (Chen et al., 2007; Zhang et al., 2003; Ye et al., 2005), and DFs have been successfully combined with other methods, including linear discriminant analysis, to indentify genes involved in pure epistasis (Zhang et al., 2003). However, this method is quite a bit more computationally expensive than RFs (Zhang et al., 2009).

Simulated evaporative cooling network analysis, tested by McKinney et al. (2009), represents an interesting blend of the ReliefF algorithm of MDR and RFs within a machine-learning evolutionary method (based upon simulating chemical reaction dynamics), which improves upon both MDR and RFs in handling large datasets involving epistasis in both simulated datasets and a new adverse reaction to small pox vaccine GWAS dataset (McKinney et al., 2009). This represents what seems to be the first attempt to combine statistical methods to overcome limitations imposed by individual methods (such as multicollinearity in RFs and search-space dimensionality in MDR), which has been urged recently by multiple experts in the field as a means to improve data analysis in genetic epidemiology and genomics (Heidema et al., 2006; Ziegler et al., 2008; Moore & Williams, 2009).

3) EVOLUTIONARY ALGORITHMS

3.1) Classical Genetic Algorithms

As datasets and analytic functions increase in complexity, nonlinearity, and size, many calculus-based optimization techniques fail, necessitating the use of enumerative techniques, such as the Expectation-Maximization algorithm or evolutionary algorithms (Tang et al., 1996; Whitley, 1994). Genetic algorithms (GAs), evolutionary strategies in computing based on the principles of evolutionary biology and population genetics created by Holland in 1975, offer quick and efficient means of solving difficult or analytically impossible problems in function optimization (such as variable selection or identification of optimal parameter weightings), ordering problems (permutation problems including the infamous Traveling Salesman Problem), and automatic programming (such as genetic programming or grammatical evolution, based off of transcription, translation, and protein folding) (Forrest, 1993; Tang et al., 1996; Harik et al., 1999; Fan et al., 2007; Wang et al., 2006; Hassan et al., 2004). Genetic algorithms, with built-in mechanisms to avoid local optima and search through very large solution spaces for global optima, thrive in situations in which other enumerative and machine-learning techniques stall or fail to converge upon global solutions (as the search space is of dimension R^N , where N represents the number of parameters in the dataset) and have been successfully employed in such fields as statistical physics (Somma et al., 2008; Ceperly & Alder, 1986), quantum chromodynamics (Temme et al., 2011), aerospace engineering (Hassan et al., 2004), molecular chemistry (Deaven & Ho, 1995; Najofi et al., 2011), spline-fitting within function estimation (Pittman, 2001), and parametric statistics (Najafi et al., 2011; Gayou et al., 2008; Broadhurst et al., 1997; Paterlini & Minerva, 2010).

GAs consist of a basic iterative framework, in which several methodological variations have been developed in each phase of the algorithm to tailor it to the problem of interest (Tang et al., 1996; Goldberg et al., 1989; Miller & Goldberg, 1996). First, an initial population of individuals representing possible solutions to a given problem, are usually generated at random (though directed evolution based upon a prior distribution is possible). These individuals, consisting of bit strings called chromosomes, encode solutions within their genes, or bits in the chromosome string, which can stand for selected variables, string length, values of parameters, or branches of computer programs. Binary alphabet representation with Gray coding of genes is generally accepted and widely used with as a gene coding mechanism, though numerical and octimal/hexidecimal alphabets exist. Populations may also be split and separated into distinct and isolated subpopulations to evolve in parallel with occasional migration of individuals between subpopulations to balance genetic drift and solution diversity; this can be accomplished by island models, mimicking evolutionary effects of systems like the Galapagos Islands, or through a cellular set up, in which individuals only interact with others in their neighborhoods and are isolated by distance (Whitley, 1994; Forrest, 1993; Whitley et al., 1998).

After an initial population (or subpopulations) is generated, individuals are evaluated, or ranked, based upon a fitness function

(such as R^2 , means square errors, BIC, least squares error, or partial least squares error in variable selection problems), which computes each solution to the original problem based on an individual's encoded genes (Tang et al., 1996; Han & Kim, 2002; Paterlini & Minerva, 2010; Najafi et al., 2011). Individuals are then selected for replication and other genetic operators based upon fitness, with more fit individuals having higher selection probabilities than less fit individuals. Selection can occur via round-robin or elimination tournament selection or by random sampling through ranking or proportional roulette selection; selective pressure is a key determinant of convergence rate, as well as an algorithm's ability to avoid local optima traps, and must be carefully chosen (Miller & Goldberg, 1996; Forrest, 1993; Tang et al., 1996; Whitley, 1994).

Once individuals are chosen, they probabilistically undergo several possible genetic operations designed to evolve the population toward a solution, including 1- or 2-point crossover (in which chromosomes pair and mate in a similar fashion to meiosis), flip or absolute mutation of chromosome bits (mimicking DNA replication errors), inversion (in which a portion of the chromosome flips its orientation), and catastrophic mutation, or "triggered hypermutation" of many individuals in a population (in the spirit of mass extinctions) upon premature convergence of a population in order to escape a locally optimal solution (Tang et al., 1996; Whitley, 1994). Restrictions on crossover, such as incest prevention, may also be employed to encourage diversity and avoid founder effects; usually, this prevents chromosomes within a certain Hamming distance (a measure of dissimilarity of bits within a pair of chromosomes) from pairing with each other for crossover (Whitley, 1991). Probabilities are assigned to each operation and affect convergence time and likelihood of finding an ideal solution (Miller & Goldberg, 1996). In a parallel model, subpopulations may also undergo global or local migration at this time. This creates a new generation of individuals.

To keep the number of individuals within a population or subpopulation constant from generation to generation, individuals are deleted after the application of genetic operators (Tang et al., 1996). Three basic methods may be instituted to accomplish such: 1) generational replacement, in which all N offspring individuals created for the new generation after undergoing genetic operation replace all N individuals of the parent generation (and risk losing an optimal solution from the older generation), 2) elitist replacement, in which the best $n\%$ of individuals in the parent generation survive and mix with the $1-n\%$ best individuals in the offspring generation (a more conservative approach than generational replacement), or 3) a steady-state mix, in which the worst n individuals of a parent generation are replaced by the best n individuals in the offspring generation (Tang et al., 1996). These individuals composing the next generation are then evaluated and selected to create the next generation; this process continues until a convergence criteria, usually until a predetermined number of generations is reached or a population evolving to within a certain restricted range of fitness value differences appears for a specified number of generations (i.e. all fitness values are within ϵ units of each other). The best of these

solutions is then selected as the solution to the problem under consideration (Forrest, 1993; Han & Kim, 2002).

In genetic algorithm theory, the algorithm does not search randomly through binary space of dimension N as it evolves a population, which would create problems for convergence in large search spaces (Forrest, 1993; Whitley, 1991); rather, the algorithm searches for good patterns within chromosomes, geometrically represented as hyperplanes within a search space (Whitley, 1991; Forrest, 1993; Goldberg et al., 1985; Nowotniak & Kutcharski, 2010). Searching through variations of these building blocks, denoted as schemas, through crossover and mutation allows the GA to identify optimal schemas and combinations of schemas quickly, while mutation also allows the GA to find a global solution by destroying locally-optimal schema to allow search for other schemas and schema combinations that may lead to a global optimum (Whitley, 1991).

GAs have been employed in statistics to solve the so-called "restrictive knapsack problem," in which a restricted number of items from a group of all items must be chosen in such a way as to optimize one of their collective properties, for instance, R^2 value in multiple regression (Han & Kim, 2002; Han & Kim, 2003; Changsheng et al., 2009; Han et al., 2001). In regression, exhaustive search of all N variables and their combinations of L items is impractical or impossible for large N or L ; however, searching through many combinations and blocks of combinations at once in a GA's evolving population (or many GA subpopulations) allows for a solution to be found in these problems (Broadhurst et al., 2011). GARST, a GA which performs this search and searches for optimal mathematic transformations of chosen variables, has shown promise in small datasets with linear and nonlinear (interaction) relationships (Paterlini & Minerva, 2010), and maybe be of use in multi-method approaches to genetic epidemiology (such as RF to filter data for a GA-optimized logistic regression model). As mentioned previously, evolutionary algorithms have been successfully combined with clustering methods and NNs to improve performance with large, complex data sets (Xiao et al., 2008; Li et al., 2001; Jirapech-Umpai et al., 2005; Motsinger-Reif et al., 2008).

3.2) Quantum-Inspired Evolutionary Algorithms

A recent development in evolutionary computing has involved borrowing principles related to quantum theory and quantum computing to reduce computational cost (in some cases exponentially) and solve problems involving larger search spaces (Rylander et al., 2001; Malossini et al., 2007). Essentially, these quantum evolutionary algorithms (QEAs) exploit superposition of states in their chromosome bits (called qubits), in which all possible states of a chromosome exist simultaneously according to each state's probability until an observation is made to collapse the system to a single chromosome of bits, and entanglement, the phenomenon of information linkage between parts of a system even when the system is separated by distance (Han et al., 2001; Han & Kim, 2002). Inference is based upon subpopulations of superposed states, and all solutions are stored at once (Rylander et al., 2001).

Rather than bits composing chromosomes, QEAs utilize qubits, which represent a mixture of bit states 0&1 with probabilities of α^2 and β^2 , respectively, depicted as:

$$x|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

where $\alpha^2+\beta^2=1$. Superposed chromosomes, then, can be represented as

$$\begin{bmatrix} \alpha_1 & \alpha_2 & \dots & \alpha_n \\ \beta_1 & \beta_2 & \dots & \beta_n \end{bmatrix}$$

(Akter & Kahn, 2010; Rylander et al., 2001). For example, the chromosome $|001\rangle$, where $\alpha^2=0.33$ and $\beta^2=0.67$, composes $2/27^{\text{th}}$ of the superposed chromosome of $n=3$. Generally, parallel initial subpopulations are created with one or more individuals within a subpopulation with $\alpha=\beta=\frac{1}{\sqrt{2}}$, suggesting an equal chance of either state for each qubit in a population's chromosomes (Akter & Kahn, 2010; Han & Kim, 2002). However, previous knowledge of the problem (i.e. expert knowledge or the use of distribution priors within a Bayesian framework) may suggest an alternate weighting of α_n and β_n to guide the algorithm to a potentially optimal solution more quickly (Han & Kim, 2002; Han & Kim, 2004).

After creating the first superposed parallel subpopulations, an observation is made to collapse the systems to binary chromosomes traditionally employed by classical GAs based on the probabilities of α and β (Han & Kim, 2002). Individuals are then evaluated and ranked according to fitness, as in GAs, and the best solution is chosen and stored as a reference; all other chromosomes undergo transformation according to a unitary operator ($UU^*=U^*U$, where U^* is the adjoint), usually a Q-gate (sometimes in conjunction with a NOT gate, which serves as a mutation operator, or replaced by a Hadamard gate), which rotates the probabilities of each qubit state toward a generational subpopulation's best solution (Han & Kim, 2002; Malossini et al., 2007). This operator, shown below,

$$U(\Delta\theta_i) = \begin{bmatrix} \cos(\Delta\theta_i) & -\sin(\Delta\theta_i) \\ \sin(\Delta\theta_i) & \cos(\Delta\theta_i) \end{bmatrix}$$

obtains its rotation angle for each qubit, $(\Delta\theta_i)$, ideally between 0.001π and 0.1π , either from a look-up truth table about the qubit of interest's relation with the best solution's qubit at that position and contribution to the problem's solution (Han & Kim, 2003) or through the use of a second evolutionary algorithm, such as particle swarm optimization (Wang et al.'s quantum swarm evolutionary algorithm, 2006). The best solution is stored, and the next generation of superposed individuals is created based upon the updated probabilities (Han & Kim, 2002). This is repeated, occasionally with an added local or global migration operator, until a convergence criterion is met to yield a global solution.

Results for complex restrictive knapsack problems are promising, and several parameter and method variations improve upon computational

cost and effectiveness in optimization. For QEA in general, migration period parameters play important roles in generating diversity and avoiding local optima; global migration every 100 generations and local migration every generation seems to provide the best balance (Han & Kim, 2002). Compared to the best classical GA (CGA) with population size of 100 evolved over 1000 generations, a QEA with a single population of 2 converged to a better solution 29 times faster than the CGA and stabilized to an acceptable solution within 30 generations (Han & Kim, 2003). Parallel QEAs (PQEAs), which include subpopulations with migration periods, outperform QEAs with much shorter run times (34 seconds QEA vs. 6 seconds for PQEA in one knapsack problem) and greater fitness values of solutions than QEAs with a single population, and both outperform classical GAs on runtimes and fitness values (Han et al., 2001).

The quantum swarm evolutionary algorithms, which use particle swarm optimization to update qubit probabilities rather than a look-up table, converge faster than QEAs when faced with large knapsacks (such as more variable combinations within large genetic epidemiology datasets) but runs more slowly than QEAs (Wang et al., 2007). For example, a knapsack with 500 items employing this algorithm consisting of a population of 30 and generation time of 1000 took about 98 seconds to converge, which is longer than a QEA with the same parameters but much quicker than other methods; however, convergence occurred in fewer generations within an excellent computational time, suggesting a convergence criteria based upon similarity of population rather than a preset number of generations (Wang et al., 2007). On function optimization problems, a similar algorithm, a hybrid QEA with PSO Q-gate update scheme (HQEA), converged to a significantly better solution than QEA or PSO (another evolutionary algorithm on its own) in less than half the number of generations and slightly less runtime than QEA (Changsheng et al., 2009).

A recent development of a modified QEA (QMEA, or quantum-inspired multiobjective evolutionary algorithm) to tackle multiobjective knapsack problems, which identifies many combinations that maximize a combined profit (such as R^2 in regression problems) within certain combinatorial restrictions (such as those imposed by MDR or K-means clustering methods), outperformed traditional methods on several knapsack problems (250 items, 500 items, and 750 items, respectively), maintaining higher diversity and higher quality of solutions over a larger search space (Kim et al., 2006). This algorithm shows promise as a wrapper search method for MDR (which could employ EVD testing to retain all significant n-way interactions found) and as a possible optimizer of clustering methods, including KNN and KMC. QEAs have already been adapted to clustering method problems, though results have varied on dataset analysis through QEA clustering of microarray data (Zhou et al., 2005); QMEA may serve as a better optimization strategy by allowing multiple objectives to be optimized and multiple solutions evolved. Work in this field has been scarce thus far.

An interesting development along the lines of using priors to weight α and β in the generation of (an) initial population(s) is the two-phase QEA (TPQEA), in which local subpopulations are isolated and evolved to a best solution in each subpopulation (without global

migration); those best solutions are then used to generate each initial subpopulation within a PQEA framework (Han & Kim, 2004). When compared to QEA performance on various restrictive knapsack problems, TPQEA converged more quickly than QEA, with time savings increasing exponentially with increases in knapsack size and item relationship complexity! More impressively, TPQEA shows nearly perfect performance on small problems with known solutions, suggesting possible use in variable selection problems (Han & Kim, 2004).

Opportunities for QEAs in genetic epidemiology abound. With their low computational cost and robust performance on complex optimization problems, QEAs could potentially improve upon the performance of existing methods utilizing GAs (such as GAKNN, GENN, GA logistic regression, and TARGET) or methods not employing GAs yet (such as MDR) and increase their ability to process large, complex data sets to yield all possible solutions (solving dimensionality problems, epistasis/plastic reaction norms, genetic heterogeneity, and multicollinearity). TPQEAs may also offer a more effective way to construct tree ensembles in a similar manner to BART with its use of estimation and optimization of multivariate priors before evolving populations to ideal solutions (as sort of a quantum TARGET/BART or quantum TARGET/RF hybrid). In addition, these methods could be combined with MDR utilizing EVD testing to identify significant n-way interactions with a data set to be entered into logistic regression or RFs with single predictors to create a model with main effects and epistatic effects. An adaptation of GARST through QEAs may also be useful in processing datasets or previously-identified subsets of variables (through RFs or QEA-KNN...) in logistic regression.

4) POTENTIAL NEW METHODS IN GENETIC EPIDEMIOLOGY WORTH CONSIDERING

4.1) Multistep Methods

The use of two or more methods has been suggested as a possible solution to the limitations imposed by single-method techniques (i.e. dimensionality and MDR, epistasis in logistic regression...). Many possible methodological combinations exist, specifically involving the use of evolutionary algorithms.

First, RF could be used to identify genetic and environmental factors associated with disease through revised importance measures for use in an MDR. Using an evolutionary algorithm (such as QMEA) or SURF and TuRF filter with an EVD test of significance would allow MDR to identify n-way interactions within the set of important predictors, which would then be fed into logistic regression with the predictors identified by RF to create a predictive, interpretable model of disease risk. If the number of predictors is too large or transformation of variables may be necessary, GARST or a quantum version of GARST could be used to optimize variable selection for the logistic regression model.

Along those lines, QMEA could first be used with EVD-based testing in MDR methods (or KNN or KMC) to identify significant n-way interactions, which could be entered into logistic regression with single predictors (with an evolutionary algorithm to reduce dimension

if the curse of dimensionality plagues the dataset). This could also involve a step using RF as a logistic regression filter for interaction and main effect terms.

Additionally, RFs on its own or in conjunction with GARST (or quantum GARST developed from one of the aforementioned QEA versions) could be used as filter for logistic regression, identifying a small subset important factors that could be tested for main effects and interaction terms.

For newly-developed methods employed in these set-ups, performance could be compared with existing methods on test/simulation datasets before use in real genetic epidemiology studies (such as comparing GARST and a quantum GARST in regression models). Multimethod results could then be compared to other methods on test/simulation data and nascent datasets to verify significant increases in computation time, model performance, and ease of interpretation through these new multistep methods.

4.2) Tree-Based Models

Several intriguing extensions of existing tree-based methods involving evolutionary algorithms exist. A quantum version of TARGET could be developed using HQEA, TPQEA, or QMEA instead of the existing GA to improve tree optimization. With increased prediction accuracy and a very simple interpretation, this method may offer a tenable alternative to hard-to-interpret ensemble techniques, such as RF or BART, or serve as a potential new basis for tree growth on subsets of variables in ensemble methods (such as the previously mentioned blend of TPQEA-optimized trees with BART or RF). These new techniques could be compared to existing techniques on UCI repository datasets or genomics datasets and then applied to new datasets if results seem promising.

4.3) Neural Network Training

Another promising possibility is the use of new and existing QEAs (such as TPQEA, QMEA, or HQEA) in the training of neural networks or optimization of neural network structure. An extension of Venayagamorthy and Singhal's multilayer perceptron NNs and simultaneous recurrent NNs could involve training with TPQEA, a faster and more accurate algorithm than other QEAs. This could then be compared to other methods, such as RF or evolutionary-algorithm-assisted logistic regression, on UCI test data sets and new real-world studies in genetic epidemiology.

4.4) Missing Data Solutions for Datasets in Genetic Epidemiology

Missing data within genetic epidemiology datasets poses statistical challenges, as existing parametric-based, explicit imputation techniques (such single and multiple imputation with Expectation-Maximization algorithms or Markov Chain Monte Carlo methods) fail when assumptions (such as the curse of dimensionality) are not satisfied (Gheyas & Smith, 2009). Implicit imputation

techniques (which don't impose many assumptions when imputing data) are few and far between, including hot or cold deck imputation (calculating missing values by evaluating similar points in space with complete data on the variable of interest), missForest (which used RFs of nonmissing predictors to compute outcomes for each missing variable), and a modified generalized regression neural network algorithm (GSI for single imputation and GMI for multiple imputation), which is based on a Euclidean distance function between points (He, 2006; Stockhoven & Buhlmann, 2011; Gheys & Smith, 2009). MissForest shows promise; however, computation time is polynomial with respect to the number of variables and longer for datasets including categorical variables (Steckhoven & Buhlmann, 2011). While reducing forest size and node split subset size effectively reduces computation time in tested datasets, it is unknown how missForest would handle very large datasets with a feasible computational cost (Steckhoven & Buhlmann, 2011). GMI offers a quick and effective imputation method compared to existing explicit methods, but its computational cost has not been reported for any size of dataset (Gheys & Smith, 2009). These techniques warrant further investigation as possible imputation methods for datasets in genetic epidemiology.

Bibliography

- Akter, S., & Khan, M. H. (2010). Multiple-Case Outlier Detection in Multiple Linear Regression Model Using Quantum-Inspired Evolutionary Algorithm. *Journal of Computers* , 1779-1788.
- Amaratunga, D., Cabrera, J., & Lee, Y.-S. (2008). Enriched Random Forests. *Bioinformatics* , 2010-2014.
- Breiman, L. (2001). Random Forests. In *Machine Learning* (pp. 5-32).
- Broadhurst, D., Goodacre, R., Jones, A., Rowland, J. J., & Kell, D. B. (1997). Genetic algorithms as a method for variable selection in multiple linear regression and partial least squares regression, with applications to pyrolysis mass spectrometry. *Analytica Chimica* , 71-86.
- Bureau, A., Dupuis, J., Falls, K., Lunetta, K. L., Hayward, B., Keith, T. P., et al. (2005). Identifying SNPs Predictive of Phenotype Using Random Forests. *Genetic Epidemiology* , 171-182.
- Bush, W. S., Dudek, S. M., & Ritchie, M. D. (2006). Parallel multifactor dimensionality reduction: a tool for the large-scale analysis of gene-gene interactions. *Bioinformatics* , 2173-2174.
- Bush, W. S., Edwards, T. L., Dudek, S. M., McKinney, B. A., & Ritchie, M. D. (2008). Alternative contingency table measures improve the power and detection of multifactor dimensionality reduction. *Bioinformatics* , 238-255.
- Ceperley, D., & Alder, B. (1986). Quantum Monte Carlo. *Science* , 555-561.
- Cha, S.-H., & Tappert, C. (2009). A Genetic Algorithm for Constructing Compact Decision Trees. *Journal of Pattern Recognition Research* , 1-13.
- Chang, J. S., Yeh, R.-F., Wiencke, J. K., Wiemels, J. L., Smirnov, I., Pico, A. R., et al. (2008). Pathway Analysis of SNPs Potentially Associated with Glioblastoma Multiforme Susceptibility Using Random Forests. *Cancer Epidemiology Biomarkers* , 1368-1373.

Changsheng, G., Juan, H., & Liang, Z. (2009). A New Hybrid Quantum Evolutionary Algorithm and Its Application. *Proceedings of the 5th WSEAS International Conference on Mathematical Biology and Ecology*, (pp. 98-102).

Chen, X., Liu, C.-T., Zhang, M., & Zhang, H. (2007). A forest-based approach to identifying gene and gene-gene interactions. *PNAS* , 19199-19203.

Chipman, H., George, E. I., & McCulloch, R. E. (2010). BART: Bayesian Additive Regression Trees. *Annals of Applied Statistics* , 266-298.

Chipman, H., Kolaczyk, E., & McCulloch, R. (1998). Bayesian CART Model Search. *Journal of the Statistical Association* , 935-960.

Clarke, J., & West, M. (2008). Bayesian Weibull tree models for survival analysis of clinico-genomic data. *Statistical Methodology* , 238-262.

Cook, N. R., Zee, R. Y., & Ridker, P. M. (2004). Tree and spline based association analysis of gene-gene interaction models for ischemic stroke. *Statistics in Medicine* , 1439-1453.

Culverhouse, R., Klein, T., & Shannon, W. (2004). Detecting Epistatic Interactions Contributing to Quantitative Traits. *Genetic Epidemiology* , 141-152.

Deaven, D. M., & Ho, K. M. (1995). Molecular Geometry Optimization with a Genetic Algorithm. *Physical Review Letters* .

Denison, D. G., Mallick, B. K., & Smith, A. F. (1998). A Bayesian CART Algorithm. *Biometrika* , 363-377.

Deutsch, J. M. (2002). Evolutionary algorithms for finding optimal gene sets in microarray prediction. *Bioinformatics* , 45-52.

Diaz-Uriarte, R., & Andres, S. A. (2006). Gene selection and classification of microarray data using random forest. *Bioinformatics* , 3-16.

Dressman, H. K., Berchunck, A., Chan, G., Zhai, J., Bild, A., Sayer, R., et al. (2007). An Integrated Genomic-Based Approach to Individualized Treatment of Patients with Advanced-Stage Ovarian Cancer. *Journal of Clinical Oncology* , 517-524.

Fan, G., & Gray, B. (2005). Regression Tree Analysis Using TARGET. *Journal of Computational and Graphical Statistics* , 1-13.

Fan, K., O'Sullivan, C., Brabazon, A., & O'Neill, M. (2007). Option Pricing Model Calibration using a Real-valued Quantum-inspired Evolutionary Algorithm. *GECCO* (pp. 1983-1989). London, England, UK: ACM.

Forrest, S. (1993). Genetic Algorithms: Principles of Natural Selection Applied to Computation. *Science* , 872-878.

Gayou, O., Das, S., Zhou, S.-M., Marks, L. B., Parda, D. S., & Miften, M. (2008). A genetic algorithm for variable selection in logistic regression analysis of radiotherapy treatment outcomes. *Medical Physics* , 5426-5433.

Gheyas, I. A., & Smith, L. S. (2009). A Novel Nonparametric Multiple Imputation Algorithm for Estimating Missing Data. *Proceedings of the World Congress on Engineering*. London, UK.

Goldberg, D. E., Korb, B., & Deb, K. (1989). Messy Genetic Algorithms: Motivation, Analysis, and First Results. *Complex Systems* , 493-530.

Greene, C. S., Penrod, N. M., Kiralis, J., & Moore, J. H. (2009). Spatially Uniform ReliefF (SURF) for computationally-efficient filtering of gene-gene interactions. *BioData Mining* , 5-14.

Hahn, L. W., Ritchie, M. D., & Moore, J. H. (2003). Multifactor dimensionality reduction softway for detecting gene-gene and gene-environment interactions. *Bioinformatics* , 376-382.

Han, K.-H., & Kim, J.-H. (2003). On Setting the Parameters of Quantum-inspired Evolutionary Algorithm for Practical Applications. *Proceedings of 2003 Congress on Evolutionary Computation*, (pp. 178-184).

Han, K.-H., & Kim, J.-H. (2002). Quantum-Inspired Evolutionary Algorithm for a Class of Combinatorial Optimization. *IEEE Transactions on Evolutionary Computing* , 580-592.

Han, K.-H., & Kim, J.-H. (2004). Quantum-Inspired Evolutionary Algorithms With a New Termination Criterion, HE Gate, and Two-Phase Scheme. *IEEE Transactions on Evolutionary Computing* , 156-169.

Han, K.-H., Park, K.-H., Lee, C.-H., & Kim, J.-H. (2001). Parallel Quantum-inspired Genetic Algorithm for Combinatorial Optimization Problem. *IEEE* , 403-406.

Harik, G. R., Lobo, F. G., & Goldberg, D. E. (1999). The Compact Genetic Algorithm. *IEEE Transactions on Evolutionary Computation* , 287-297.

Hassan, R., Cohanin, B., de Weck, O., & Venter, G. (2004). A Comparison of Particle Swarm Optimization and The Genetic Algorithm. *Jet Propulsion* , 1-13.

Heidema, G. A., Boer, J. M., Nagelkerke, N., Mariman, E. C., van der A, D. L., & Feskens, E. J. (2006). The challenge for genetic epidemiologists: how to anlayze large numbers of SNPs in relation to complex disease. *Genetics* , 23-38.

Hothorn, T., Lausen, B., Benner, A., & Radespiel-Troger, M. (2004). Bagging Survival Trees. *Statistics in Medicine* , 77-91.

Hua, J., Xiong, Z., Lowey, J., Suh, E., & Dougherty, E. R. (2005). Optimal number of features as a function of smaple size for various classification rules. *Bioinformatics* , 1509-1515.

Jiang, R., Tang, W., Wu, X., & Fu, W. (2009). A random forest approach to the detection of epistatic interactions in case-control studies. *The 7th Asia Pacific Bioinformatics Conference*, (pp. 565-577). Beijing, China.

Jirapech-Umpai, T., & Aitken, S. (2002). Feature selection and classification for microarray data analysis: Evolutionary methods for identifying predictive genes. *Bioinformatics* , 48-59.

Kim, Y., Kim, J.-H., & Han, K.-H. (2006). Quantum-inspired Multiobjective Evolutionary Algorithm for Multiobjective 0/1 Knapsack Problems. *2006 IEEE Congress on Evolutionary Computation* (pp. 9151-9156). Vancouver, BC, Canada: IEEE.

Klein, R. J. (2007). Power analysis for genome-wide assoication studies. *Genetics* , 58-66.

Lee, J. W., Lee, J. B., Park, M., & Song, S. H. (2005). An extensive comparison of recent classification tools applied to microarray data. *Computational Statistics and Data Analysis* , 869-885.

Lee, S. Y., Chung, Y., Elston, R. C., Kim, Y., & Park, T. (2007). Log-linear model-based multifactor dimensionality reduction method to detect gene-gene interactions. *Bioinformatics* , 2589-2595.

Li, L., Weinberg, C. R., Darden, T. A., & Pedersen, L. G. (2001). Gene selection for sample classification based on gene expression data: a study of sensitivity to choice of parameters of the GA/KNN method. *Bioinformatics* , 1131-1142.

Loh, W.-Y., & Shih, Y.-S. (1997). Split Selection Methods for Classification Trees. *Statistica Sinica* , 815-840.

Lou, X.-Y., Chen, G.-B., Yan, L., Ma, J. Z., Zhu, J., Elston, R. C., et al. (2007). A Generalized Combinatorial Approach for Detecting Gene-by-Gene and Gene-by-Environment Interactions with Application to Nicotine Dependence. *The American Journal of Human Genetics* , 1125-1136.

Lunetta, K. L., Hayward, B., Segal, J., & Van Eerdewegh, P. (2004). Screening large-scale association study data: exploiting interactions using random forests. *Genetics* , 32-45.

Malossini, A., Blanzieri, E., & Calarco, T. (2007). Quantum Genetic Optimization. *IEEE* , 1-30.

Maulik, U., & Bandyopadhyay, S. (2000). Genetic algorithm-based clustering technique. *Pattern Recognition* , 1455-1465.

McKinney, B. A., Crowe, J. J., Guo, J., & Tian, D. (2009). Capturing the Spectrum of Interaction Effects in Genetic Association Studies by Simulated Evaporative Cooling Network Analysis. *PLoS Genetics* , 1-12.

Meng, Y. A., Yu, Y., Cupples, A., Farrer, L. A., & Lunetta, K. L. (2009). Performance of random forest when SNPs are in linkage disequilibrium. *Bioinformatics* , 78-95.

Miller, B., & Goldberg, D. L. (1996). Genetic Algorithms, Selection Schemes, and the Varying Effects of Noise. *Evolutionary Computation* , 113-133.

Moore, J. H., & Williams, S. M. (2009). Epistasis and Its Implications for Personal Genomics. *American Journal of Human Genetics* , 309-317.

Moore, J. H., Asselbergs, F. W., & Williams, S. M. (2010). Bioinformatics Challenges for Genome-Wide Association Studies. *Bioinformatics* , 445-455.

Motsinger-Reif, A. A., Dudek, S. M., Hahn, L. W., & Ritchie, M. D. (2008). Comparison of Approaches for Machine-Learning Optimization of Neural Networks for Detecting Gene-Gene Interactions in Genetic Epidemiology. *Genetic Epidemiology* , 325-340.

Najafi, A., Ardakani, S. S., & Marjani, M. (2011). Quantitative Structure-Activity Relationship Analysis of the Anticonvulsant Activity of Some Benzylacetamides Based on Genetic Algorithm-Based Multiple Linear Regression. *Tropical Journal of Pharmaceutical Research* , 483-490.

Nowotniak, R., & Kucharski, J. (2010). Building Block Propagation in Quantum-Inspired Genetic Algorithms. *Automatics* .

Ooi, C. H., & Tan, P. (2002). Genetic algorithms applied to multi-class prediction for the analysis of gene expression data. *Bioinformatics* , 37-44.

Paterlini, S., & Minerva, T. (2010). Regression Model Selection Using Genetic Algorithms. *Recent Advances in Neural Networks, Fuzzy Systems, and Evolutionary Computing* , 19-26.

Pattin, K. A., White, B. C., Barney, N., Gui, J., Nelson, H. H., Kelsey, K. R., et al. (2009). A Computationally Efficient Hypothesis Testing Method for Epistasis Analysis using Multifactor Dimensionality Reduction. *Genetic Epidemiology* , 87-94.

Pittman, J., & Murthy, C. A. (2001). Fitting optimal piecewise linear functions using genetic algorithms . *IEEE Transactions on Pattern Analysis and Machine Learning* , 701-718.

Pittman, J., Huang, E., Dressman, H., Horng, C.-F., Cheng, S. H., Ysou, M.-H., et al. (2004). Integrated modeling of clinical and gene expression information for personalized prediction of disease outcomes. *PNAS* , 8431-8436.

Pittman, J., Huang, E., Nevins, J., Wang, Q., & West, M. (2004). Bayesian analysis of binary prediction tree models for retrospectively sampled outcomes. *Biostatistics* , 1-15.

Qi, Y. (2011/2012). Random Forest For Bioinformatics. In *Ensemble Learning: Methods and Applications*.

Robison, A. J., & Nestler, E. J. (2011). Transcriptional and epigenetic mechanisms of addiction. *Nature Reviews Neuroscience* , 623-635.

Rylander, B., Soule, T., Foster, J., & Alves-Foss, J. (2001). Quantum Genetic Algorithms. *Proceedings of the Genetic and Evolutionary Computing*, (pp. 1005-1011).

Segal, M. R. (2003). Machine Learning Benchmarks and Random Forest Regression. *Center for Bioinformatics and Molecular Statistics* .

Sexton, R. S., Dorsey, R. E., & Johnson, J. D. (1998). Toward global optimization of neural networks: a comparison of the genetic algorithm and backpropagation. 1-36.

Somma, R. D., Boixo, S., Barnum, H., & Knill, E. (2008). Quantum Simulations of Classical Annealing Processes. *Physics Review Letters* , Letter 101.

Stekhoven, D. J., & Buhlmann, P. (2011). MissForest--nonparametric missing value imputation for mixed-type data. *Oxford Journal's Bioinformatics* , 1-12.

Strobl, C., Boulesteix, A.-L., Zeileis, A., & Hothorn, T. (2007). Bias in random forest variable importance measures: Illustrations, sources and a solution. *Bioinformatics* , 25-46.

Tang, K. S., Man, K. F., & He, Q. (1996). Genetic Algorithms and their Applications. *IEEE Signal Processing Magazine* , 22-36.

Temme, K., Osborne, T. J., Vollbrecht, K. G., Poulin, D., & Verstraete, F. (2011). Quantum Metropolis Sampling. *Nature* , 87-90.

Venayagamoorthy, G. K., & Singhal, G. (2005). Quantum-Inspired Evolutionary Algorithms and Binary Particle Swarm Optimization for Training MLP and SRN Neural Networks. *Journal of Computational and Theoretical Nanoscience* , 561-568.

Wang, Y., Feng, X.-Y., Huang, Y.-X., Pu, D.-B., Zhou, W.-G., Liang, Y.-C., et al. (2007). A novel quantum swarm evolutionary algorithm and its applications. *Neurocomputing* , 633-640.

Whitley, D. (1994). A Genetic Algorithm Tutorial. *Statistics and Computing* , 65-85.

Xiao, J., Yan, Y., Lin, Y., Yan, L., & Zhang, J. (2008). A Quantum-inspired Genetic Algorithm for Data Clustering. *IEEE* , 1513-1518.

Ye, Y., Zhong, X., & Zhang, H. (2004). A genome-wide tree- and forest-based association analysis of comorbidity of alcoholism and smoking. *Genetics* , S135-140.

Zhang, H., Wang, M., & Chen, X. (2009). Willows: a memory efficient tree and forest construction package. *Bioinformatics* , 130-135.

Zhang, H., Yu, C.-Y., & Singer, B. (2003). Cell and tumor classification using gene expression data: Construction of forests. *PNAS* , 4168-4172.

Zhou, Z.-H., Wu, J.-X., Jiang, Y., & Chen, S.-F. (2001). Genetic Algorithm based Selective Neural Network Ensemble. *Proceedings of the 17th International Joint Conference on Artificial Intelligence* (pp. 797-802). Morgan Kaufmann.

Ziegler, A., Konig, I. R., & Thompson, J. R. (2008). Biostatistical Aspects of Genome-Wide Association Studies. *Biometrical Journal* , 8-28.