

# NUANCE: Naturalistic University of Alberta Nonlinear Correlation Explorer

GEOFF HOLLIS and CHRIS WESTBURY  
*University of Alberta, Edmonton, Alberta, Canada*

In this article, we describe the Naturalistic University of Alberta Nonlinear Correlation Explorer (NUANCE), a computer program for data exploration and analysis. NUANCE is specialized for finding nonlinear relations between any number of predictors and a dependent value to be predicted. It searches the space of possible relations between the predictors and the dependent value by using natural selection to evolve equations that maximize the correlation between their output and the dependent value. In this article, we introduce the program, describe how to use it, and provide illustrative examples. NUANCE is written in Java, which runs on most computer platforms. We have contributed NUANCE to the archival Web site of the Psychonomic Society ([www.psychonomic.org/archive](http://www.psychonomic.org/archive)), from which it may be freely downloaded.

“Facts are nothing without their nuance, sir.”  
—Norman Mailer, upon being told by the judge to stick to the facts during the 1969 trial of the Chicago Seven

Genetic programming (GP) is a way of using natural selection to program computers (Koza, 1992). Unlike genetic algorithms and neural networks, which use idiosyncratic and human-unfriendly binary representations of problems and their solutions, GP evolves ordinary computer code. The output it produces are strings of code in a human and machine comprehensible language. Westbury, Buchanan, Sanderson, Rhemtulla, and Phillips (2003) suggested that this makes GP particularly useful in solving the kinds of nonlinear multivariate regression problems that are common in psychology and other scientific disciplines and discussed some of the statistical and conceptual issues involved in using such tools in scientific work. In solving nonlinear multivariate regression problems, it is often highly desirable to be able to see exactly how multiple independent variables relate to a particular dependent variable. In this article, we introduce a free new tool that has been specially designed for this purpose: NUANCE (Naturalistic University of Alberta Nonlinear Correlation Explorer). Written in Java, NUANCE will run on most computer platforms. It can explore and suggest human-readable solutions for any problems in which one or more predictors may predict the value of a single dependent value. It is easy to use, fast, flexible, and powerful, in-

corporating a new rank-order-based fitness function that is demonstrably superior to the averaged multitest fitness (described below) introduced by Westbury et al.

Our goals in this article are twofold. First, we aim to describe the interface of NUANCE, so that anyone will be able to use it after reading this article. Second, we aim to demonstrate that NUANCE works and that the novel algorithm it introduces is superior to averaged multitest fitness, which was itself introduced to address problems of particular relevance to curve fitting.

## How Does GP Work?

The general principles of GP are simple. In this section, we describe them in general terms; in later sections, we will fill in the details as they apply to NUANCE.

To begin, one needs to define two elements: a *function set* and a *fitness value*.

The function set consists of the set of computational operators that are available for use in any evolved solutions. Although these may be highly specialized for the specific problem under consideration, they may also consist simply of a set of plausibly relevant basic mathematical functions.

The fitness value is a way of quantifying the goodness of any evolved solution. These may be a very general measure, such as a squared error, a percent hit rate, or a correlation of an evolved function's output with known correct values. The function that calculates the goodness of any solution is called the *fitness function*.

Once a fitness function and a function set are defined, evolution of solutions may begin. In the first generation, many (at least hundreds, more often thousands) of potential solutions are randomly generated, by conjoining the operators in the function set in legal ways. Since these solutions are randomly generated, they are likely to be very poor at solving the problem. However, because the fitness function is defined, these thousands of poor solutions can

---

This work was made possible by a National Engineering and Science Research Council grant from the Government of Canada to C. Westbury. The authors gratefully acknowledge the many useful suggestions received in response to an earlier version of this article and the software from Jonathan Vaughan, Greg Francis, Ian Walker, an anonymous reviewer, and Gail Moroschan. Correspondence should be addressed to C. Westbury, Department of Psychology, P-220 Biological Sciences Bldg., University of Alberta, Edmonton, AB, T6G 2E9 Canada (e-mail: [chrisw@ualberta.ca](mailto:chrisw@ualberta.ca)).

be rank ordered for their utility in solving the problem. The best functions are identified, by using some selection criterion, and retained. The rest of the population is discarded. The selected functions are randomly broken into subparts of random size, ranging from a single node to the whole tree. The subparts are then randomly conjoined to fill the population of the next generation.

In sum, there are three steps to GP: generate functions, select the best, and regenerate new functions by mixing up the selected best functions. As this brief description suggests, the process used in GP is closely analogous to selective breeding in biology. In selective breeding, the breeder decides which animal is suitable to breed, on the basis of what characteristics are desired. In GP, the breeder's choice is automated, instantiated in the fitness function and selection criterion. The animals to be bred are mathematical equations.

By repeated application of the three simple steps, GP can incrementally improve the average and best fitness of each succeeding generation (for formal analyses, see Holland, 1992, and Koza, 1992). The process can be stopped when a "perfect" solution is found, if such a solution exists, or it may be stopped because some other completion criterion is met: for example, because the slope of the fitness function is close to 0, or because some set number of generations has passed.

Multivariate estimation problems of the type handled by NUANCE are well suited for GP because they have a natural fitness function. In such problems, fitness may be simply defined as the absolute value of the correlation of the estimate with the data to be estimated. Because correlations are well defined and continuous across the bounded range of  $-1$  to  $1$ , it is possible to compare the output of any two estimator equations and have the computer unambiguously decide which one is better: the one with the larger absolute correlation with the data to be estimated.

This brief overview should make clear that understanding any GP system is largely a matter of understanding three things: the function set, the fitness function, and the selection criteria. In the first section that follows, we consider NUANCE's function set. In the second section, we consider its fitness function and selection criteria, which are combined in NUANCE.

## UNDERSTANDING NUANCE

### Function Set

NUANCE uses a flexible approach to its function set. It has a set of functions built in, which users may turn on or off. NUANCE also allows users to add their own custom functions.

Most of the built-in functions are simple well-known functions of one or two inputs: addition, subtraction, multiplication, division, natural logarithm, square, cube, square root, cube root, absolute value, sine, cosine, and tangent. All these functions are enabled by default.

NUANCE does not include a general power function, since it is too easy to generate overflow errors using

large (absolute) powers. However, note that all low-power and small-root functions can easily be constructed by NUANCE during evolution by concatenating the square, cube, square root, and cube root functions. Moreover, such power functions may be explicitly added to the user-defined functions.

NUANCE also includes four multi-input built-in functions. The first is an if/then/else statement. This statement takes three arguments. If the first argument is 1, it returns the second argument; otherwise, it returns the third argument. The remaining three functions are the ">" (greater than), "<" (less than), and "=" (equals) comparators. They each take two arguments, returning 1 if the specified relationship holds, and 0 otherwise. The use of these functions allows NUANCE to evolve radically nonlinear solutions, which have discontinuities. However, their use also causes huge increases in the size of evolved solution trees, thereby slowing down the evolutionary process and decreasing the probability that the solution will be humanly comprehensible. Moreover, it is also possible to use nested if/then statements to generate trivial solutions to any problem, by building lookup tables that mirror the input file. These functions should therefore be used judiciously: Include them whenever there is reason to suspect that the solution may be very complex or discontinuous, and exclude them by turning them off if not. These functions are off by default.

It is our intention to define all the functions that most users will want inside NUANCE, and we will add more as they are deemed necessary. However, we have also provided a means for users to add their own functions. At the bottom of NUANCE's Function Set panel, there is a button called Custom Functions. When clicked, users can add and delete their own programmable functions. Users are required to provide a unique name for each function and a description of what the function does. Descriptions must be written in prefix notation, with the operator in front of its arguments, just like all functions NUANCE produces. The input parameters for each function must be designated within the function description using special symbols (N1, N2, N3, . . . , Nn). When the function is used, subtrees attached to the custom function will replace these symbols.

For example, the function (+ N1 2) will add two to the subtrees attached to it. The function (\* N1 (\* N1 N1)) will return the cube of any node it takes as an input value. It is allowable to use more than one parameter in a user-defined function. For example, the function (+ (\* N1 (\* N1 N1)) (sin N2)) will add the sine of one input to the cube of another.

In general, simple combinations of primitive built-in functions that are useful to the problem need not be explicitly defined by the user, since they may be reasonably expected to be discovered by the evolutionary process itself. In addition, it is possible that the provision of very fit custom functions may harm the evolution process. If the fitness provided by a custom function is much greater than that provided by the basic operators, NUANCE may have a tendency to converge on a solution that is nothing

more than the custom function, rather than exploring the solution space as it is intended to do. Custom functions may not necessarily be useful and should be defined only if necessary.

### Fitness Function and Selection Criteria

In this section, we discuss NUANCE's fitness function, which ranks evolved equations, and its selection criterion, by which a subset of ranked equations is selected for breeding to fill the next generation.

Any solution to a problem that generalizes from a small set of data to a general relation faces the danger of overfitting, by finding solutions that are highly suited to the data set on which they were evolved, but that generalize poorly to other data sets. We say that such solutions are *brittle*. Brittleness is defined as an evolved solution's inability to generalize to unseen problem cases that are part of the problem of interest but that were not used as a fitness set for the GP system. For instance, if one were to try to evolve a control behavior to direct the foraging of an artificial ant, an ant's behavior would be considered brittle if it could not forage with competence in an environment different from the one it is trained in. For NUANCE to have any utility, it must be able to create generalizable summaries of the data sets that it models. In other words, NUANCE must avoid overfitting its solutions or producing brittle solutions.

Kushchu (2002) and Moore and Garcia (1997) suggest the same solution for avoiding brittle solutions: the use of randomized learning cases. If we change the problem that is used to assess fitness on each generation, it presents an environment that selects for generalization. For example, in a scenario of an evolved simulated ant, Kushchu did this by randomizing the starting location of the ants on each generation. He showed that this produced much more adaptive behavior to new environments than did having the ants start their foraging from the same place each generation.

In NUANCE, we utilize an analogous strategy, by using a new random subset of the input data to gauge the fitness of each generation. One problem with this method is that it runs the risk of culling out evolved functions that actually do generalize. This can happen if a subset were generated that happened by chance to be highly unrepresentative of the problem of interest. Poor fitness with that subset could wipe out many evolved functions that actually do generalize, thus setting back the evolutionary process.

Westbury et al. (2003) introduced a technique called *averaged multitest fitness* (AMTF), which lessened, though it did not necessarily remove, this problem of losing fit functions. AMTF builds cross-validation into its fitness function, in a way that is analogous to the  $k$ -fold cross-validation technique (Stone, 1974) that is often used in training neural networks. The algorithm takes the average correlation of an equation's output with the measure to be predicted over multiple random subsets. It uses this average as a gauge of an equation's fitness.

AMTF has the disadvantage of requiring multiple fitness tests for every evolved function. This makes it com-

putationally costly, thereby slowing down the speed at which evolution can proceed.

We have recently developed a novel fitness method that is superior to AMTF for a number of reasons, which we call *age-weighted fitness*. This is the default method used within NUANCE, which we discuss in the Runs section below.

For further discussion of the problem of overfitting, see Example 3 below.

## USING NUANCE

NUANCE is simple to run. In this section, we describe the steps.

NUANCE is written in Java. Java is freely available for most common computer operating systems. NUANCE requires Java 1.4.2 or later. If you need a copy of the Java Virtual Machine in which NUANCE runs, you can download one for most platforms ([java.sun.com/j2se/downloads.html](http://java.sun.com/j2se/downloads.html)).

NUANCE can be downloaded from the Psychonomic Society archive ([www.psychonomic.org/archive](http://www.psychonomic.org/archive)). NUANCE comes packaged as a Java archive (JAR) file and is run by double-clicking on the file icon. Java archives open more slowly than regular applications, so be patient.

After opening NUANCE, you will see a graphical interface with three tabs labeled "NUANCE," "OUTPUT," and "FORMAT CONVERTER." We will describe each of these three panels in turn.

### NUANCE Panel

The NUANCE panel is the main panel for using the program. On this panel, you can control the main function of NUANCE by setting the evolution parameter values and the input and output files, choosing which functions to allow and turning the evolutionary process on and off.

In the pane to the left, you will see check boxes for each of the built-in functions described above. You can turn these on or off by clicking on these buttons.

The middle pane on the right, entitled "Input and Output Files," allows you to specify the input files and output directories for NUANCE. The input file must be a tab-delimited text file. By default, the dataset.txt file goes in the same directory as NUANCE. You may change this by clicking on the "change" button beside the file path field or by typing a new file path directly into that field. The first line of the input file must consist of the variable names, separated by tabs. Beneath this are the columns of data, also separated by tabs. The first column must contain the independent measure to be predicted. The remaining columns must contain predictors. NUANCE does not accept missing data; therefore, missing values must be replaced with an estimate, or rows that contain missing values must be deleted.

You must also specify a directory where NUANCE will print its output. This defaults to a time-stamped directory, whose name reflects the time NUANCE was started. This directory will be written by default inside the "logs" directory at the same level as NUANCE itself.

In each output directory, NUANCE will create subdirectories that are numbered for each run: Run 1, Run 2, and so on. Each of these output subdirectories contains two files: `history.txt` and `log.txt`. The `history.txt` file will contain a log of all the output NUANCE displays on screen. The `log.txt` file contains information about the whole run that might be of interest: average time per generation, average length and age of functions in that generation, the best function of that generation, and its age and fitness.

When a specified set of runs is completed, each output directory also contains four other files: `settings.txt`, `functions.txt`, `best_estimates.txt`, and `graphing_data.txt`.

The `settings.txt` file simply records the parameter settings that were used for those runs.

The `functions.txt` file contains a list of all functions, including custom functions, used during the runs.

The `best_estimates.txt` file contains the best evolved function across all runs, in both prefix and infix notation, as well as a list of the output of that function for each data point in the input file. It is possible that some cases may be undefined. This can happen if a particular case that was never selected as part of the testing subset of a given function is undefined for that function. Undefined values are replaced with “NaN” (i.e., “not a number”).

The `graphing_data.txt` file contains a different representation of the same best function. It includes one  $10 \times 10$  table for each pair of predictors. Each table spans across the range of the two predictors in 10 equal steps. Each of the 100 cells in the  $10 \times 10$  table contains the normalized value of the output function over the 10 steps of the two predictors. If there are more than two predictors in the data file, the others are fixed at their average value for the purposes of building the table. The results in each cell are normalized (expressed in  $z$  scores) because the output units of evolved functions are arbitrary and therefore may be difficult to interpret. The purpose of this table is to facilitate visual inspection of the shape of the best function and of the nature of the relation between any two predictors. The output table can be pasted into any other program with graphing capabilities in order to create the graph. We have included with NUANCE an Excel template (`GraphingTemplate.xls`) for creating such graphs.

NUANCE’s leftmost pane lists of all functions and operators used by NUANCE. The default setting is to turn all functions on, except for the *logical if* and the *greater, less than, and equal* operators. This default can be overridden by putting a list of the function names you wish to use the `functions.txt` file located in NUANCE’s directory. Custom functions can also be added to this list; the format for doing so is to place the function’s name and description on a single line, separated by a tab.

In the top pane on the right, you will see a set of all NUANCE’s parameters and their values. NUANCE will always load its default parameters from the `settings.txt` file (if the file exists) located in the same directory as NUANCE. You can edit this file directly to change the default settings that NUANCE loads with, or you can click on the Save Settings button to write the displayed settings as the default settings for future use. You can also change

the settings manually in the top right-hand pane. We now describe the parameters in detail.

### NUANCE’s Population Parameters

**Population Size.** The Population Size parameter specifies how many equations can be in the evolving population at any time. In general, this variable can be fruitfully maximized, but with two caveats: (1) the time it takes to run with larger populations will increase linearly with population size, and (2) NUANCE may require more RAM. The default heap size for Java applications is 64 MB. If memory needs to be increased, run the program from the command line. To do so, type “`java -mx###m -jar nuance.jar`,” where `###` is the number of megabytes of RAM you would like to allocate to NUANCE. One hundred twenty-eight megabytes of RAM should be sufficient for all populations under size 10,000.

The default population size is 2,500, which is probably fine for most problems. Increase it if you doubt that the best solution is being found. A population size above 10,000 would be very large indeed. We have successfully run NUANCE with populations as large as 20,000.

**Generations.** The Generations parameter specifies how many generations you wish to evolve. It is difficult to give a general value to be entered here, since the number of generations required may vary with the complexity of the problem. As with many of the other NUANCE parameters, the rule of thumb is the more, the better. Since, in the end, you will probably keep only the best solution, you can never do worse by running more generations, and you will probably (and, in the long run, almost certainly) do better. However, running more generations takes more time.

It is possible to estimate if one has run sufficient generations by looking at the shape of the fitness curve, which graphs fitness by generation (for examples, see Figures 1, 2, 5, and 6). If the curve is nearly flat, then little progress is being made in each generation, and it is unlikely that further computation will lead to substantial improvement. If the curve has a clear positive slope, more computation will probably lead to further gains. However, this heuristic may sometimes be misleading. The stochastic nature of evolutionary search means that it is possible for progress to be made after many generations of no progress. An example of this phenomenon is shown in Figure 1, which shows a NUANCE run, using a public data set, that tried to predict average January temperature in 60 cities in the United States from average rainfall in those cities. After 47 generations with an almost flat fitness function ( $r = .11$ ), the run graphed here found a substantially better region of search space, and in one generation jumped to  $r = .29$ . Thereafter, fitness steadily increased. For this reason, it is advisable to allow leeway in the number of generations of search. The default value is 100 generations.

**Runs.** You can set the Runs parameter to any  $N > 0$ . NUANCE will repeat evolution  $N$  times, with all of the settings already defined. It is desirable to repeat runs several times because of the stochastic nature of evolutionary computing. There is no guarantee that any run is the best possible. We recommend that you run at least 8–10 times,



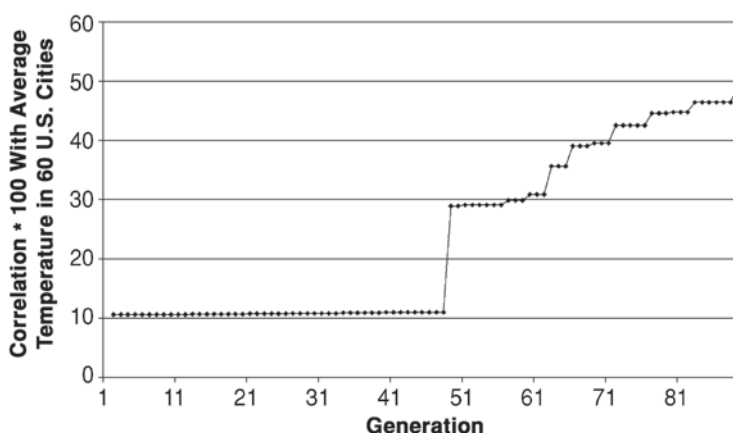


Figure 1. A NUANCE fitness curve, using Median Pass selection, illustrating the danger of over-reliance on the shape of the fitness curve as a measure of whether or not one has run sufficient generations. This problem used a data set from DASL, at [lib.stat.cmu.edu/DASL/Datafiles/SMSA.html](http://lib.stat.cmu.edu/DASL/Datafiles/SMSA.html), to estimate the average January temperature in 60 U.S. cities from average rainfall in those cities. The fitness curve is almost perfectly flat with a correlation of approximately .11 for 47 generations, which might lead one to conclude that no better solution was possible. However, a breakthrough at generation 48 led to substantial and thereafter increasing improvement on this early estimate. This illustrates why it is a good idea to run as many generations as possible.

or more if you can, although clearly the best solution in hand may be deemed sufficient for any number of reasons at any given time. The default value is 10.

**Starting Tree Depth.** Users are allowed to specify how large their starting functions will be by specifying how deep each function tree is. The default starting depth is 3, which means an average of 5 or 6 nodes per function. The final functions produced by NUANCE rarely exceed depths of 10, and this parameter should not be set higher than that. Such deep functions are very difficult for humans to understand.

Although the default starting tree depth should suffice for most purposes, longer and shorter values may be of benefit in some situations. Without allowing shorter starting trees, it is possible that you may miss some very simple solutions. However, this is not a very serious problem, because you will almost certainly be able to deduce the missed solutions from the solutions that NUANCE offers (see Example 4C below). In some cases, it may be desirable to set the starting depth higher than the default. It is possible, when the best solution tree is deep, that short solutions may lead to premature convergence to regions of solution space that are suboptimal. If you encounter a problem in which you feel dissatisfied that a good solution has been found, you can try increasing this parameter.

**Parsimony Pressure.** One of GP's major limitations is that functions may get so large that they are completely incomprehensible, intractable to run, or a combination of the two. A number of theories have been proposed to explain this phenomenon, which is commonly termed *bloat* (see Langdon, 2000; Soule & Foster, 1999; Streeter, 2003). Some (Soule & Foster, 1998; Streeter, 2003) speculate

that large size is actually evolutionarily adaptive, since it can lessen the probability that crossover will damage offspring. Others (e.g., Langdon & Poli, 1997) suggest that bloat may be caused by distribution of fit functions in solution space. Solutions with equal fitness can have many different physical forms. With variable-length functions and any arbitrary cutoff to separate long from short, there are typically more long solutions than short solutions; thus, long solutions have a bias to be found first.

Whatever the cause of bloat, it is undesirable from the NUANCE user's point of view. We have implemented parsimony pressure to prevent functions from growing excessively large. Parsimony pressure penalizes long solutions. By setting parsimony pressure to a positive value, a function's fitness is decreased by a percent value equal to the parsimony pressure times the number of nodes (operators and arguments) in the function. For example, a parsimony pressure of 1 would decrease the fitness of a function whose length is 100 by 1%. For a length of 150, fitness would decrease by 1.5%, for a length of 200, fitness would decrease by 2%, and so forth.

A small parsimony pressure can have a large effect on function size, since it discourages increases in function length except where a long function is markedly better than the best short function. On the data set that is discussed as Example 1 below, even the smallest parsimony pressure of 1% reduces average function length to less than one fifth of its value without such pressure (see Figure 2). The default parsimony pressure for NUANCE is 1.

**Mutation Rate.** Randomly mutating members of an evolving population is a way to add diversity to the population pool. Many researchers working with GP have

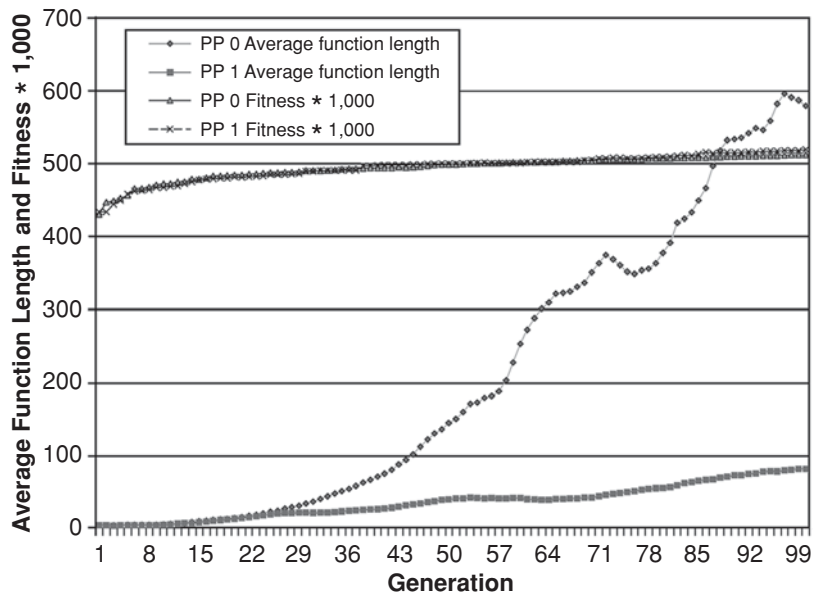


Figure 2. Average function length and fitness (\*1,000) by generation, with (PP 1) and without (PP 0) the minimal parsimony pressure of 1%. This figure shows that even a small parsimony pressure (penalty for long functions) can have a marked effect on average function length, while having little effect on fitness. In order to show both fitness and length on the same graph, we have multiplied fitness values (correlations with the predictor) by 1,000. These data are from the ON × FREQ problem discussed in Example 1.

deemed mutation to be generally ineffective (Koza, 1992; Luke & Spector, 1998). However, it may be a worthwhile addition to NUANCE in some specific situations. If many solutions become too similar too quickly (i.e., they converge to a local optima), evolution may halt. Mutation can help prevent such problems by introducing new genes into the population pool. It does so by randomly changing some percentage of nodes in the children that are seeded into a population. The default mutation rate is 0, which means that there will be no mutation. Users may wish to increase this value if the best function’s fitness stops increasing very early in the run. If fitness still fails to increase even with mutation on, it is probable that there are very few solutions better than the best found. If, on the other hand, turning on mutation helps to increase fitness, it is likely that there was a problem with getting stuck on local optima.

**Testing Subsets and Subset Size.** As mentioned in the introduction to this article, in order to avoid problems with overfitting, a GP program can use small portions of the data set for testing population members. Overall fitness is the average of the fitness from all subsets. In this way, the GP system treats the input data set as if it were made up of multiple smaller subsets and tries to find solutions that generalize between these subsets. In NUANCE, there are two parameters associated with this ability: Testing Subsets and Subset Size.

The first parameter, Testing Subsets, is the number of different subsets that the population members are tested on

each generation. In some cases, any particular subset may be unrepresentative of the problem being solved, which means that there is always a chance of killing fit functions because they cannot accurately predict unrepresentative data. Increasing the value of this parameter to some number greater than 1 allows NUANCE to use the averaged multitest fitness described above. Each evolved function is tested on multiple subsets. Averaging the function’s fitness on all subsets gives the function’s fitness for the generation. Data will never be included more than once within any given subset. However, they may be included more than once across multiple subsets.

The second parameter, Subset Size, is the size of the data set(s) that will be used for testing the population members on each generation, expressed as a percentage of the total data set size.

We consider these two parameters together here because their optimal values depend on each other. The general rule of thumb is that the Testing Subsets parameter should be as high as possible (so each function is tested on as many subsets as possible), and the Subset Size parameter should be as low as possible (so that each subset is small relative to the overall data set size). However, these optimal settings depend to some extent on how many rows your input set has. A Testing Subsets parameter value above 2 is desirable in order to maximize generalizability, but it is better to have at least 15 or 20 rows in each subset than it is to have many subsets. If the Subset Size is so low that there are very few data points in each subset, no

evolution will occur: The system will simply thrash about randomly, unable to find an equation with a consistent fit. Finally, the number of rows specified by combining these two variables should ideally be a relatively small fraction of the total data set size. If you specify a Testing Subsets size of 2 and a Subset Size of 50%, you will be using your entire input data set to assess the fitness of every function, which will expose you to the risk of overfitting.

The default values are a Testing Subsets size of 1 and a Subset Size of 20%, on the assumption that large data sets are rare. These parameters mean that the fitness of each function is tested on a random fifth of the input data set. If you have more than 100 lines of input, you should increase Testing Subsets to 2 and perhaps decrease Subset Size to 15%. This will assess every function on two random subsets of the input data set, each composed of 15% of that data set. Every function will therefore be tested on a randomly selected 30% of the input data set.

**Fitness Age Weight and Selection Strategy.** Large data sets are rare. In problems where one cannot collect large samples of data, it is impossible to use multiple subsets to prevent unrepresentative data from killing fit population members without overfitting. NUANCE therefore includes other tools to deal with this problem. By adding a bonus to fitness based on a function's age (the number of generations for which it has existed), NUANCE can get a measure of how well functions fare when dealing with different data. If a very unrepresentative testing subset is generated, this age weighting can act as a protective buffer for fit functions if they have survived for multiple generations (presumably because they are, indeed, fit).

The Fitness Age Weight parameter specifies how strong the weighting should be. The exact way that age weighting affects fitness is as follows:

$$\begin{aligned} <\text{new fitness}> = <\text{old fitness}> \\ & * (1 + \log<\text{fitness age weight}>(\text{age})). \end{aligned}$$

The benefits of age weighting are dependent on the Selection Strategy parameter—that is, how the program decides which functions should be allowed to produce offspring at the end of each generation. NUANCE offers two selection strategies, which can be chosen by a pull-down menu: Greedy Overselection and Median Pass selection. Age weighting is useful with the second strategy but cannot help with the first. In order to understand why, it is necessary to understand how each strategy works.

Greedy Overselection was the selection method introduced by Koza (1992) and used by Westbury et al. (2003). It is usually called *fitness-proportionate selection with greedy overselection*. Under this selection scheme, the population pool that represents 16% of the normalized fitness is selected for reproduction with an 80% probability, and the rest of the population is selected for reproduction with a 20% probability. Koza (1992, p. 98) specifically recommended it for one third of difficult problems in his book that required populations of more than 1,000. However, computer technology has advanced greatly since 1992. We can think of no reason to ever use a population

smaller than 1,000, and we suggest using populations several times that size.

Greedy Overselection does not work with age weighting because it allows too small a proportion of the population to survive each generation. It thereby eliminates most of the important age information that age weighting needs. Because of this limitation, we have implemented Median Pass selection. This selection scheme removes the least fit 50% of the population in each generation. It replaces that 50% with children of the top 50%, in such a way as to guarantee that every member of that population will get to parent at least one child. The second parent of each child is chosen at random from the parenting population.

NUANCE automatically implements a redundancy check with both selection schemes. This replaces any identical twins in the population with new functions, in order to maximize the population diversity and minimize computational redundancy.

NUANCE's default settings use Median Pass selection, because it is the more computationally efficient selection strategy. There are two reasons for this. One is that Median Pass selection is not so dependent on (though it may still benefit from) having a large data set that allows averaging of multiple measures for each fitness value, because it can spread out that averaging across generations. The other is that the fitness barrier imposed by Median Pass selection is more conservative than that imposed by Greedy Overselection. The use of age weighting imposes a penalty on new equations that, by chance, just happen to do well on the current subset of the input file that is being used for assessing fitness. In contrast, when using Greedy Overselection as a selection strategy, an equation that does well because it happens to be especially fit to the current fitness subset will be allowed to seed the next generation with its children. Median Pass selection is also more likely to promote functions that were created earlier, because age counts. Since older functions are almost always shorter than newer functions, equations tend to get longer and more complex under Greedy Overselection than they do under Median Pass selection. A related result is that evolution under Greedy Overselection is usually slower than evolution under Median Pass selection.

Greedy Overselection has one advantage over Median Pass selection: It tends to perform slightly better at accounting for variance in the dependent variable, for the same reason that it is slower (i.e., because larger, more complex equations can evolve and be tested). The prices to be paid for this increase in explained variance are two-fold: (1) slower evolution and (2) much more difficulty in understanding the evolved equations. Whether or not this trade is worthwhile will depend on the nature of the problem under consideration.

We have conducted several test runs to try to determine what the best value for the Fitness Age Weight parameter should be. Our tests have made clear that large bases for the age-logging function tend to be better than low bases, which overvalue age relative to fitness. However, it has proven difficult to definitively select one large base

value over another, since there is much variance in performance on different data sets. In the end, NUANCE uses  $(\log_{1000}(\text{age}+1) + 1)$  as the default value.

One piece of motivating evidence for this value is shown in Figure 3. In this example, we used NUANCE to predict lexical decision times using a large set of lexical variables as predictors. Figure 3 graphs the probability that a run using a given age-weighting factor would contain the most predictive equation in each of three consecutive bins of 33 generations. The probability was maximized in the final bin of 33 generations when the weighting scheme used a base of 1,000. Moreover, only this weighting scheme showed a monotonically increasing probability of having the best function across all three bins.

Note that Figure 3 also clearly shows that lower base values perform worse. Bases below 500 have a much lower probability of containing the best function relative to higher bases.

Notwithstanding this evidence, Figure 4 illustrates why we need not worry too much about whether or not we have identified the optimal value for the age-weighting scheme. This figure graphs the average correlation (on the same data set as used in Figure 3) of the best-performing function in 20-generation bins for several different age-weighting schemes. The final performance of the different schemes is closely comparable. The correlations of the output of best functions in the final generation (Generation 100) with the value to be predicted range between .5399 (using  $\log_{5000}$ ) and .5484 (using  $\log_{2000}$ ), amounting to a negligible difference of 0.06% in the amount of

variance accounted for in the dependent variable by the best and worst evolved equations at Generation 100.

**Local Search Priority.** GP is designed to be a global search technique. It sends multiple agents out into the search space, all in their own directions, each trying to get to the best point in the fitness landscape around them, rather than each population member attempting to get to the best-known point in the fitness landscape. Because of this, it may be the case that the best-known places in the search space are not as rigorously explored as they could be.

The Local Search Priority parameter allows the user to force “micro-exploration” around the best-known solution. It does so by generating a number of mutant offspring of the best function in the population pool. The number of mutants to be created is equal to the value of the Local Search Priority parameter. Each mutant is created by replacing one random subtree of the best function with a new, randomly generated subtree. The fitness of each mutant offspring is compared with that of its parent. If an offspring’s fitness exceeds that of its parent, it is added to the population pool. Otherwise, it is discarded. This may potentially bring the total population size to Local Search Priority +  $N$ , although this will happen only if all mutant offspring exceed the fitness of their parent. Population size never increases beyond this size, since each new generation is reset to the specified population size before any mutants are added.

We have not yet been able to document any compelling evidence that searching this space with mutation does in fact provide an evolutionary advantage. However, since

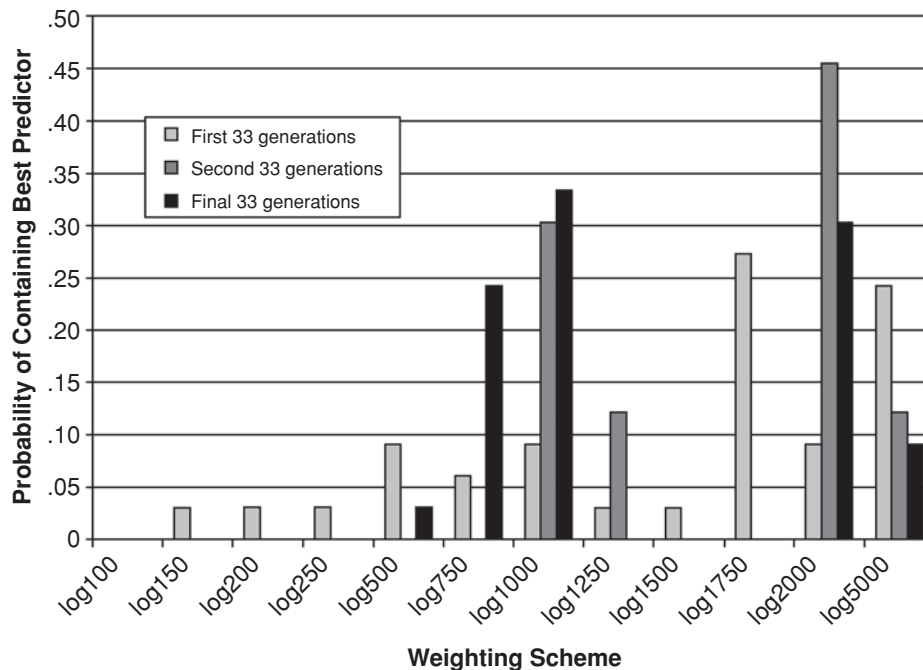


Figure 3. Probability of containing best predictor, by generation and age-weight function (averaged across bins of 33 generations). The dependent variable in this example is lexical decision RT using 10 predictors.



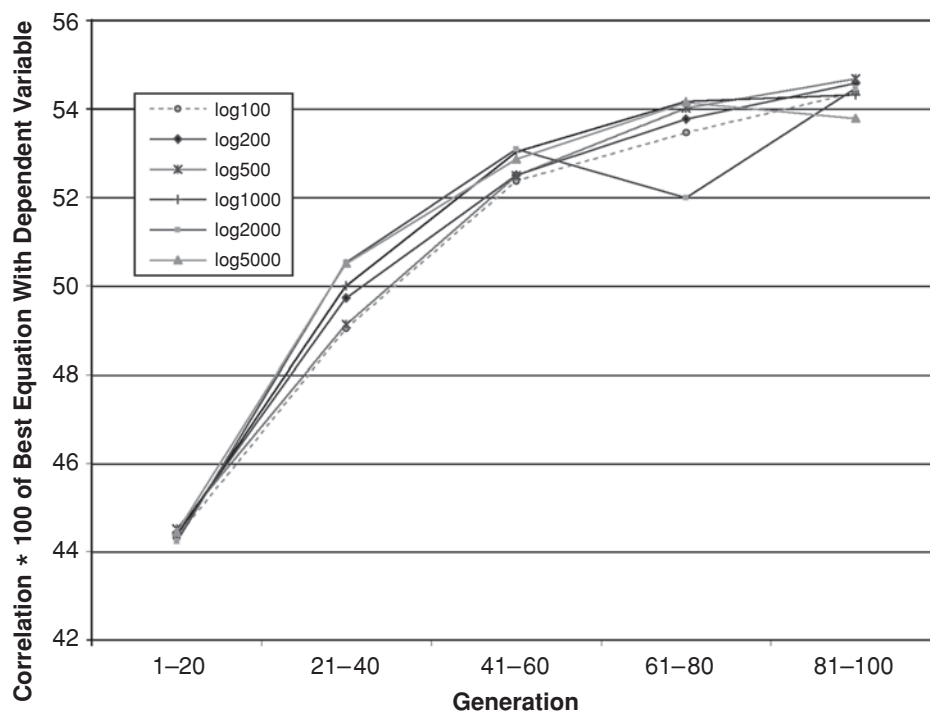


Figure 4. Correlation of the best evolved predictor equation with a dependent variable, in 20-generation bins, by age-weighting scheme. The dependent variable in this example is lexical decision RTs using 10 predictors.

the cost of allowing mutation of the best function is low and the potential advantages seem clear, we have left it in as an option.

It is possible that using too large a value for this parameter may lead to exponential growth in the population of mutants that are related to the “ancestral” best solution, with a resultant underexploration of other fruitful regions of the solution space. We therefore suggest that this parameter be set low relative to the total population size. The default value is 20.

**Functions to Display.** The Functions to Display parameter allows you to specify how many of the best functions from each generation you want to display. If the value is 1, only the single best function of each generation will be shown. The optimal value for this display parameter is purely a matter of personal taste. The default value is 5.

### OUTPUT Panel

The second tab in the graphical user interface takes the user to the OUTPUT panel. This is a tool intended for testing the goodness of fit and generating graphable tables of any evolved function on any data set. It can be used to examine functions from NUANCE’s evolutionary history, to cross-validate functions from one data set to another, and to explore how small changes may impact on a function’s fitness.

There are two fields. The Function field, on top, is for the function to be examined—it can be pasted or typed in. The second, the Dataset field, allows the user to specify

with which data set the equation in the Function field should be used. When the user clicks on the “print output” button, the specified equation is applied to the specified data set. The output contains the first column of the input file (the dependent variables) and the value generated by applying the specified equation to each row of that input file. Beneath these two columns, their correlation is given. Finally,  $10 \times 10$  tables are printed for each pair of predictors. These are the same as the tables, described above, that are written to the graphing-data.txt file: Each table ranges across the range of the two predictors in 10 equal steps, and each cell contains the normalized value of the output function over the 10 steps of the two predictors. When there are more than two predictors, the others are fixed at their average value.

### FORMAT CONVERTER Panel

The final panel is the FORMAT CONVERTER panel. This contains tools for re-representing NUANCE’s prefix equations, which place the operator at the front of its arguments. This representation is uncommon and may be particularly difficult to decode when the equations are complex because of deeply nested parentheses. The converter is easy to use: Simply paste or type a prefix equation into the top panel and then click on one of the three converter buttons to convert to a new notational system. Three notational systems are available. The first is prefix notation, the more commonly used but less consistent representation that places the operator between its arguments, when

there is more than one argument, and before its argument, when there is just one argument. The second representation is that used by the mathematical software, Mathematica (Wolfram Research, 2003). This notation can be pasted directly into a Mathematica notebook, in order to simplify the equations using Mathematica’s “simplify” command. Finally, prefix equations can be converted to LaTeX notation. This is a notation that can be pasted into compatible programs to produce typeset mathematical output. Many LaTeX typesetting programs are available for free.

**FOUR EXAMPLES**

In this section, we briefly present four examples of how NUANCE can be used.

**Example 1: Predicting Lexical Decision RTs From Two Lexical Variables**

Our first example uses the same data set as was used in Example 1 of Westbury et al. (2003), which looks at how well orthographic neighborhood (ON) and word frequency can predict lexical decision RTs. This data set was originally chosen because the relationship between the variables has been well studied and is of theoretical interest and because data are easy to obtain. It is of additional interest here because it provides us an opportunity to test NUANCE against a similar (but more limited) GP system, the Common Lisp based on the Greedy Overselection system described by Westbury et al.

ON is a measure of how many words in a given language are exactly one letter different from a target word. For example, the neighborhood of the word *free* includes

the words *tree*, *flee*, and *fret*, among others. Decades of experimental work (see Andrews, 1997, for a review) have demonstrated that ON has an effect on lexical decision task RTs. In a lexical decision task, subjects have to decide as quickly as possible whether a letter string displayed on the screen is a word or a nonword. Decisions to high-ON words are made more quickly than those to low-ON words when the words are low frequency. ON has no effect on high-frequency words. Westbury et al. (2003) showed that GP was able to deduce this relationship with a high degree of precision by fitting the two parameters to lexical decision RTs from a large arbitrary set of words. Using eight 75-generation runs of size 2,500, they evolved equations of the two variables whose output correlated with 450 RTs at .48 ( $p < .01$ ). The equation generalized to 150 unseen data points with a correlation of .61 ( $p < .01$ ). Relying on linear regression equations would cause one to underestimate the size of the effect. Linear fits correlated with the input data set with  $r = .22$  ( $p < .01$ ) and generalized to the unseen data set with  $r = .20$  ( $p < .01$ ).

We used NUANCE to estimate the same 450-item data set as had been used by Westbury et al. (2003), using both Greedy Overselection and Median Pass selection. We ran 10 runs of 100 generations of population size 5,000. Sample fitness curves for this problem are shown in Figure 5.

The best evolved predictor correlated with the RTs in the fitness set at  $|r| = .52$  using Greedy Overselection ( $p < .0001$ ) and  $|r| = .48$  using Median Pass selection ( $p < .0001$ ), about the same as the best previously found solution. The two estimates correlated with each other with  $r = .87$  ( $p < .0001$ ). They correlated very highly ( $p < .0001$ ) with the best estimates from Westbury et al.

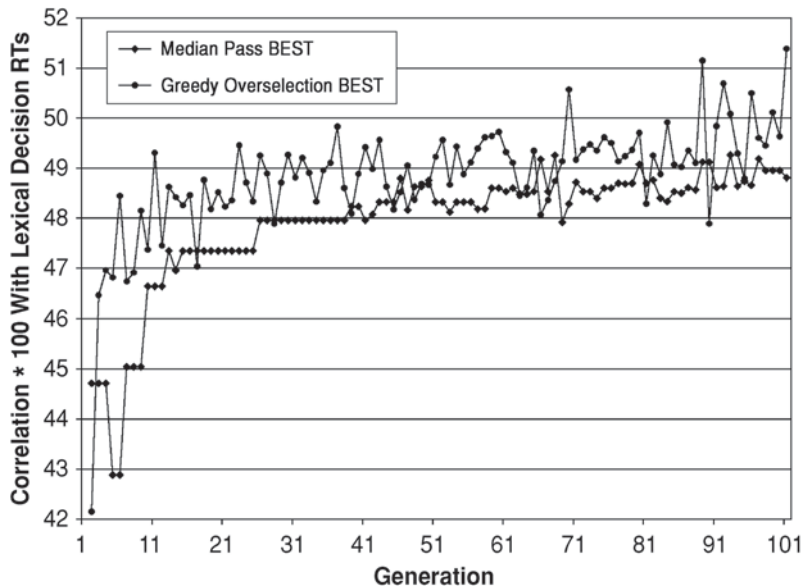


Figure 5. Two typical evolutionary runs of the best evolved predictor, by generation, for the two selection schemes offered in NUANCE. Median Pass selection with age weighting is always less erratic, but Greedy Overselection (which cannot benefit from age weighting) almost always performs slightly better at predicting variance. The dependent variable in this example is lexical decision RTs using two predictors (see Example 1 in the text).

(2003), at  $|r| = .86$  (Greedy Overselection solution) and  $|r| = .99$  (Median Pass selection solution).

In order to examine the extent to which the two evolved predictor equations can generalize to unseen data, we tested their ability to predict the RTs for 150 words that had not been in the fitness set. The best evolved predictor equation using Greedy Overselection correlated with the RTs in the entire fitness set at  $|r| = .52$  ( $p < .0001$ ). The best evolved predictor equation using Median Pass selection correlated with the RTs in the entire fitness set at  $|r| = .58$  ( $p < .0001$ ). These two predictions correlated with each other at  $r = .75$  ( $p < .0001$ ). The Westbury et al. (2003) equation produced output that correlated with those unseen data at  $.60$  ( $p < .0001$ ). The correlation of the present output with that earlier estimate is  $.98$  ( $p < .0001$ ) for the estimates derived by Median Pass selection and  $.74$  ( $p < .0001$ ) for the estimates derived by Greedy Overselection.

The best evolved equations are shown in their raw form in the Appendix, which illustrates the difference in complexity of the two solutions. The equation evolved using Greedy Overselection is far more complex than the elegant simple equation evolved with Median Pass selection. The complex solution derived with Greedy Overselection was the best one at predicting RTs in the original input set. However, it had obvious weaknesses: It was highly complex, and it correlated poorly with RTs from a related data set and with other solutions evolved on the same data set. The simpler Median Pass selection equation is almost as good at predicting RTs in the input data set as the complex equation, but it also generalizes very well to the new data set.

Overall, the fitness and test set correlations produced by NUANCE are all significantly correlated with those reported previously and discussed above, using the same data set but an entirely different GP system. The close convergence of three different evolved estimates using

different computers, fitness functions, and methods suggests that we may have a high degree of confidence that the solutions discovered are indeed the best solutions that may be found using GP.

### Example 2: Estimating Leukemia Incidence From Per Capita Cigarette Consumption

Our second example uses published data from the Data and Story Library (DASL, at [lib.stat.cmu.edu/DASL/](http://lib.stat.cmu.edu/DASL/)). Originally published in Fraumeni (1968), this data set presents the relationship between average cigarette consumption per capita and leukemia deaths per 100,000 people in 44 states in the U.S. in 1960. The linear correlation between cigarette consumption and leukemia deaths in this data set is insignificant, at  $-.07$  ( $p > .05$ ). A cubic regression yielded the best (but also statistically unreliable) correlation we could find using the standard regression options in SPSS, with  $r = .32$  ( $p > .05$ ). We used NUANCE to see whether it could find a transformation of the cigarette consumption predictor that would yield a better prediction of leukemia deaths from cigarette consumption.

We ran 10 runs of 100 generations of population size 2,500, using the default settings for other values. The average and best runs are shown in Figure 6. The best equation that was evolved was a complex equation that correlated with leukemia deaths at  $.53$  ( $p < .001$ ). This suggests that cigarette consumption may have a stronger relationship to leukemia deaths than originally suggested, though the relationship may be a complex one.

The NUANCE-evolved estimate and the raw leukemia data are graphed, in standardized units, against cigarette consumption in Figure 7. This graph suggests hypotheses that may shed light on the relationship under consideration. In particular, the idealized NUANCE estimate sug-

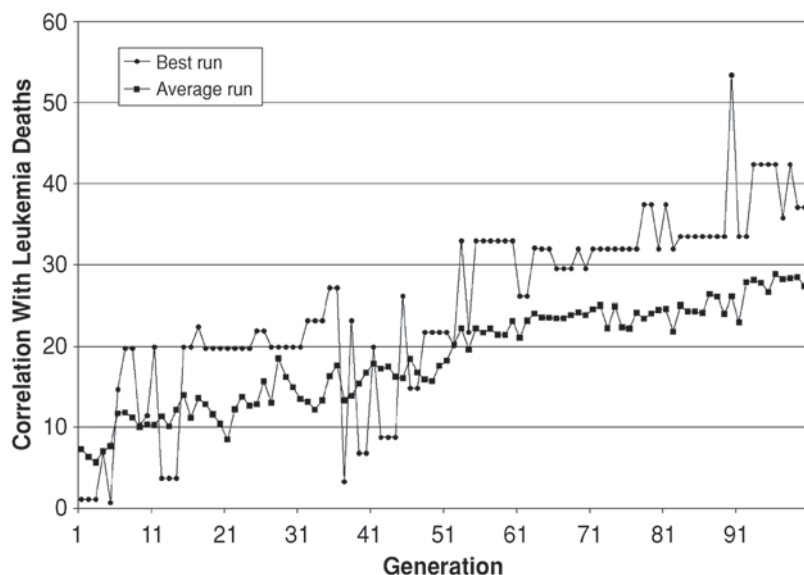


Figure 6. Best and average of the 10 NUANCE runs for Example 2, predicting leukemia deaths from cigarette consumption.

gests that the relationship between leukemia incidence and cigarette consumption may be quite linear when the number of cigarettes is low (below 23, which is the turning point of the sideways parabolic shape of the NUANCE estimates in the graph). Above this number of cigarettes, the correlation no longer holds and may even be reversed.

This possibility can be easily verified by going back to the original data. For the 21 states with average per capita consumption of at most 23 cigarettes, there is a statistically reliable linear correlation between leukemia incidence and cigarette consumption, with  $r = .50$  ( $p < .05$ ). For the 23 states with average per capita consumption above 23 cigarettes, the correlation between leukemia incidence and cigarette consumption is  $-.06$  ( $p > .05$ ). The suggestion drawn from a brief glance at NUANCE's idealization of the data holds.

The proper interpretation of this fact is outside of our own range of knowledge. Perhaps leukemia incidence is lower when cigarette smoking is higher because the smokers are dying of other smoking-related disorders before they contract leukemia, or perhaps there are other mediating factors for which cigarette consumption is an appropriate proxy variable. For the purposes of this article, we merely wish to demonstrate that NUANCE was able to quickly create a simplified summary of a complex data set, which shed light on the structure of the relationship between variables and suggested hypotheses that may be worthy of closer study.

**Example 3: The Significance of Random Data**

One criticism that is often raised against curve-fitting programs such as NUANCE is that they generate spuri-

ous solutions by overfitting data. The third example is intended to both illustrate this danger and to discuss how the danger may be avoided and appropriately conceptualized when using NUANCE.

We used Excel's random number generator to generate 500 quadruplets of random numbers. We then ran NUANCE to find a function that used the last three numbers to predict the first. We ran 10 runs, with a population size of 10,000, using the default settings with Median Pass selection, with all primitive functions enabled. The best solution found used two of the three predictors to produce an equation that correlated with the first column at  $r = .17$ , which is a highly reliable correlation ( $p < .001$ ) for 500 items.

Since the input was random, this highly reliable correlation may seem dismaying. The fact that curve-fitting programs can reliably fit random data is what moves critics to dismiss their solutions as "spurious" or "overfit." However, the terminology used in such dismissals is misleading as to the real nature of the problem. If NUANCE returns an equation that fits random data reliably, it can only do so because the relation described by the equation actually exists. The algorithm used in NUANCE does not allow it to return erroneous fits. The equations returned cannot be spurious or overfit (as least with respect to the input data set alone; see the discussion below); however, they certainly can be misleading or uninteresting. The temptation to use the stronger derogatory terminology springs from a confusion of *description* with *explanation*. NUANCE can offer a mathematically guaranteed description of data relations, but it can offer no help at all in explaining those relations. It is up to human beings to decide

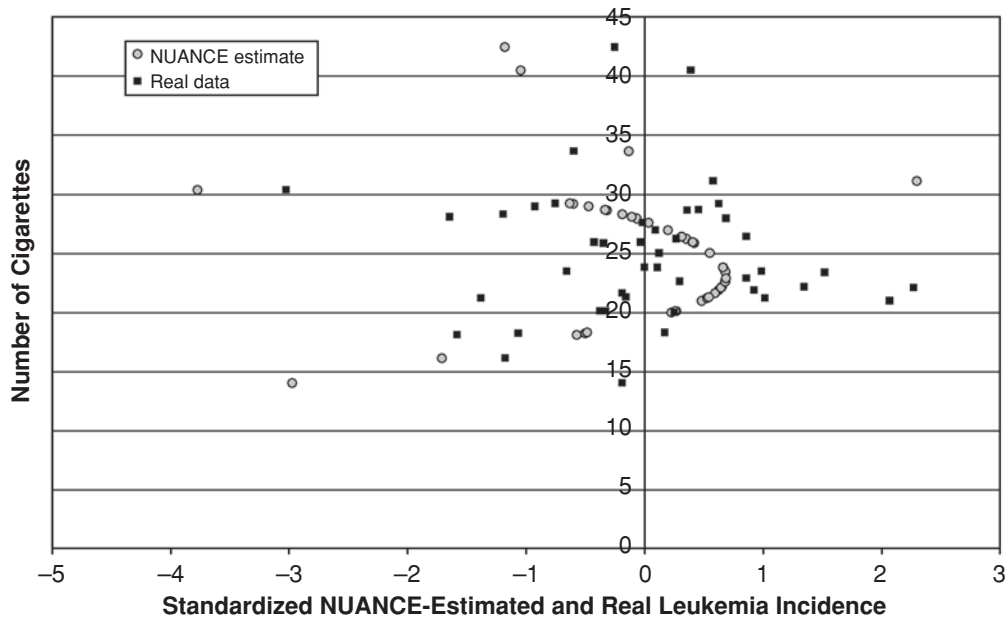


Figure 7. Predicting leukemia deaths from cigarette consumption. The NUANCE-evolved estimate and the raw leukemia data are both graphed, in standardized units, against cigarette consumption.



whether any particular description of a data set is of any theoretical or practical utility.

Westbury et al. (2003) discussed this problem and its relation to degrees of freedom and psychological theoretical concepts in some detail in the concluding discussion of their article introducing GP as a tool for psychological theorizing. Their conclusion was that

GP is a way of cloning one's imagination. By providing it with hypothetical constructs and operators, one supplies the prior understanding that constrains the search for an appropriate representation of relationships in the data. The fitness function instantiates and provides a quantifiable measure of what constitutes a solution worth imagining.

Equating a search through relation space to imagining ways of representing a problem is not a rejection of the role of theory or of the hypothetico-deductive method in science. It is simply a recognition of the fact that theoretical claims and empirically based descriptions of the relations between theoretical elements must be separate. Theories can postulate which hypothetical constructs may play a role in any particular phenomena and may certainly attempt to specify the precise nature of that relation. However, when it comes to questions of the real nature of the relation between those constructs, theoretical claims cannot trump empirical facts. If a systematic relation between hypothetical constructs can be reliably demonstrated, by any means, theory is not in a position to rule this empirical finding in or out; it can only adapt to the finding, as it must adapt to any other relevant empirical evidence. (p. 215)

The question of practical importance is: How can we decide whether any evolved solution is "worth imagining" or not? There can be no general answer to this question, because what counts as a worthwhile relation must depend largely on the current state of knowledge with respect to the issues under examination, the goodness of fit of the solution, and other specifics of each case. However, Westbury et al. (2003) offered two heuristics for assessing the worthiness of solutions evolved with GP.

The first suggestion is to avoid giving the method inputs that might have a spurious relation to the problem—that is, to avoid using GP for wild fishing expeditions (Westbury et al., 2003, p. 215). If we restrict ourselves to using NUANCE to describe relationships in which we have a preexisting interest, then the question of how to interpret reliable random relationships simply does not arise.

If we choose to ignore this heuristic by using NUANCE to search for relationships in which we have no a priori interest (and we admit to having given in to this temptation ourselves), then it is incumbent on us to temper the conclusions drawn from such an undertaking and to seek independent sources of confirmation for any conclusions we may wish to draw. When used for fishing expeditions, curve-fitting programs such as NUANCE must be only a starting point, not an end point.

The second suggestion can help us to temper our conclusions appropriately. This is the suggestion (following Stone, 1974) that we cross-validate any curve-fitted solutions by holding back some of the data that we give as

input to our curve-fitting program and then testing the best predictor equation on this unseen data set. NUANCE has within-data-set cross-validation built into its selection mechanism; it judges the fitness of evolved solutions on random subsets of the input data set. However, it can still find solutions that capture data-set-specific variance, thereby overestimating the general goodness of fit of the evolved solution. What it really means for a solution to be overfit is precisely that the solution relies on data-set-specific variance. Data-set-specific solutions will by definition perform very badly when tested on a data set different from the one from which they were derived. To illustrate this, we took the best-fitting equation to our random data set above and used it to predict a different random data set of 600 items. The fit was not reliable ( $r = -.009, p > .05$ ), clearly demonstrating what we already knew: There is no general solution to the problem of fitting randomly generated data.

Together, these two suggestions make clear why we need not be worried to discover that there are statistically reliable relations in a random data set. Since the relation is of no theoretical or practical interest, and since it fails to generalize to a related data set, we have no reason to waste any calories in pondering it.

#### Example 4: Four Simple Problems

The three examples above suggest that NUANCE can fit data in complex ways. Does NUANCE choke when the fit is not complex? Is GP too powerful for its own good, obscuring the possibility of a simple equation through brute computational force? The answer to both questions is no. If there is a simple, shallow equation that is a very good description of the fit, NUANCE usually finds it very quickly. Indeed, it offers an embarrassment of riches, by representing the solution in various different ways. The four simple examples in this section illustrate this for several sets of perfect and noisy data. All were run on an 867-MHz Power PC G4 Macintosh with 640 MB under System 10.3, while a variety of other processes were running.

The first example used the equation

$$x = \sin(\text{sqr}(a)). \quad (4A)$$

We built a data set of 500 pairs (with random values for  $a$  between 0 and 1) that fit this equation and presented the problem to NUANCE, with a population size of 10,000 and AWF to assess fitness. In the first generation (i.e., by random search, before any evolution has occurred), NUANCE found the perfectly correlated solution:

$$((\sin a) * (\cos a)) [= .5 * \sin(\text{sqr}(a))].$$

It also reported the trivially equivalent solutions  $(\sin(a * a))$  and  $\sin(\text{sqr}(a))$  in that generation. Since a perfect solution had been attained, the experiment was halted at this point, having taken 40 sec.

We repeated the same experiment with noise added to  $x$ , so that the correlation between  $\sin(\text{sqr}(a))$  and the value of  $x$  was .929. NUANCE returned all three of the solutions above in the first generation, after 29 sec. We let the

program continue for 10 runs of 100 generations. The best fit after that time was

$$(\sin ((\sin(a)/\cos (6)) * (\cos a))),$$

or

$$(\sin ((\sin(a)/0.96) * (\cos a))).$$

This correlated with the dependent variable at  $r = .932$ , adding a minuscule correction factor to the target function  $\sin(\text{sqr}(a))$ , with which it correlated at .9999.

The second example used two predictors  $a$  and  $b$ , related by the representationally shallow (if conceptually deep) equation made famous by Einstein:

$$x = a * \text{sqr}(b). \tag{4B}$$

We used the same NUANCE parameters as above.

When the data were perfectly clean, the following exact solution was returned in Generation 1, after 8 sec, before evolution had begun:

$$((\text{sqr } b)/(1 a)) [= 1/a/ \text{sqr}(b), \text{ or } a * \text{sqr}(b)].$$

We added noise to the equation, so that  $x$  correlated with the target function at .818. After 9 sec of computing, all five of the top five equations returned from Generation 1 were perfectly correlated with this solution: one stating it precisely, and the other four merely including an additional irrelevant constant. At the end of 100 generations, the best solution offered was a more complex equation that correlated with the target variable only slightly better, at .821. That equation correlated with the “canonical” solution at .999, suggesting that it was indeed picking up on the same variation.

As its name suggests, NUANCE was developed mainly to aid in finding the shape of nonlinear relations between variables. Excellent tools already exist for characterizing the parameters of relationships that are definitely known to be linear. Such tools are better suited for characterizing those linear relationships than NUANCE, since they are designed specifically to return the solution in its canonical linear form, which NUANCE is not. Nevertheless, NUANCE is able to recognize and characterize linear relationships as easily as any other relationships. Moreover, in doing so, it may open our eyes to the new ways of construing those linear relationships.

The final two examples we consider here are randomly chosen linear equations with two variables. We continued to use the same NUANCE parameters as in the above two examples.

The first equation was

$$x = 1.3a - 144.2b + 163.5. \tag{4C}$$

We began again with non-noisy data, with both  $a$  and  $b$  given random values between 1 and 100. The best equation to appear from the first generation, which took 14 sec to run, was

$$(- (/ 1.25 a) (- 5.46 b))$$

or, in prefix notation,  $1.25/a - 5.46b$ . Within the range of  $a$  (1–100), the value of  $1.25/a$  is likely to be very small on

average, so this equation already suggests what we know to be true from the target equation: The role of  $a$  must be small or nonexistent.

By Generation 22, after a few minutes of computation, the best equation correlated with the entire set at .999999, which is probably good enough for anything except a test with a known perfect answer.

The following perfect solution evolved at Generation 65:

$$(+ (\sin 8.77) (- (+ (* b (\text{sqr } 10.23)) (- (* b 6.08) (\text{cbrt } 1.09))) a)).$$

Ignoring the constant values that have no effect on the goodness of a correlational fit, this is easily simplified to

$$a - 110.92b.$$

This is the original equation divided by 1.3.

We added noise so that the correlation between the canonical output of the equation and the dependent measure given to NUANCE was .711. After 100 generations, the best equation correlated with the entire test set slightly better, at  $|r| = .724$ . In order to assess whether this solution was the correct one, we correlated the evolved solution with the 500 real values of the predictor equation. The correlation was very high ( $|r| = .97$ ), indicating that the evolved solution was indeed functionally similar to the target equation.

The best evolved equation was complex. Rather than simplifying it as above, we will demonstrate an alternative way of approaching the simplification process when using NUANCE. This method uses the evolved solution and NUANCE’s normal output to find a better model. Having the best solution in hand sets a bound that enables us to easily search for a simpler model of our phenomenon that is more easily comprehensible and almost as good. In NUANCE’s log file, we found (among several other candidates worth considering) the very simple function  $b * \ln(b)$ , which NUANCE reported to be correlated with the predictor variable almost as well as the final best evolved equation:  $|r| = .70$ . The difference in the variance accounted for by these two models—that is, from the best evolved equation versus the extremely simple equation,  $x = b * \ln(b)$ —is less than 1%.

If this were a real problem, we could surmise from this finding that predictor  $a$  is not a relevant factor in predicting the dependent variable. In this toy example, since we have the true values of the phenomenon being modeled, we are able to test whether  $b * \ln(b)$  is a good model. The values of  $b * \ln(b)$  correlate with the output of the equation to be predicted at  $|r| = .997$ . Since we have a mathematically precise definition of the phenomenon under study, we are also able to establish whether or not this is a good general solution, by again conducting Monte Carlo simulations as we did above. Over 20 random sets of 500 pairs, the average absolute correlation between  $b * \ln(b)$  and the output of the original linear equation defined over two variables is .997 ( $SD = .0001$ ). Knowing the equation as we do, it is not surprising that this is such a good model. In fact  $b$  alone, which is very highly correlated with  $b * \ln(b)$ , is an even better model. NUANCE

is unlikely to find single-node solutions, so it missed the simpler and better model  $b$ . However, it was nevertheless able to quickly guide us to see that it is only variable  $b$  that really matters in this example.

In the final example, we look at a less skewed linear relation with integral coefficients:

$$x = 5a + 8b + 6. \quad (4D)$$

We again used random values between 1 and 100 for  $a$  and  $b$ . The best equation from the first generation was

$$(+ (/ a 10) (\text{sqrt } b))$$

or, in infix notation,  $a/10 + \text{sqrt}(b)$ . This correlated with the 500 target values at .982.

A perfect solution was found at Generation 28 of the first run and expressed succinctly as

$$(+ (+ (* b 4) a) (+ a (/ a 2))).$$

We simplified by inspection to

$$2.5a + 4b.$$

This is the original equation divided by 2.

We added noise so that the actual dependent values correlated at .672 with the target values. The best function over 100 runs correlated with those target values only slightly better, at .677. The correlation between the output of this evolved equation and the values given by the original equation was .988, suggesting, as the previous noisy examples did, that the best solution captures the same variance as the original target equation.

We simplified the best function by analogy, as above, by finding a simpler function in NUANCE's log files that did almost as well as the best equation. The following equation correlated with the target values at .676:

$$(+ 2.44 (+ (+ (+ b a) (+ a (+ (+ b b) b))) (- (+ b a) (\tan a))))).$$

This equation was selected because it can be easily simplified by inspection (ignoring the constant) to

$$3a + 5b - \tan(a).$$

This equation is a good model of this data set, since it produces output that correlates with the actual values of the original target equation in the input at .991. However, it is somewhat specialized to that particular file: A Monte Carlo simulation over 20 sets of 500 sets of two random numbers related by the target equation shows that its average correlation with the original target equation is .70 ( $SD = .09$ ). However, since the file itself contained values that correlated at only .68 with that target equation, we may not reasonably expect an exact reconstruction of that target equation.

Together, these four final examples illustrate two main points. The first is that NUANCE is always able to find the exact solution to fit data that have a known simple characterization. This gives us confidence in the solutions it offers to real problems, where the characterization of

the problem is not known in advance. The second point is that the evolutionary history reported by NUANCE can be a fruitful source of hypotheses, potential models, and general "mathematical commentary" on the phenomenon under study. By looking for equations that are both attractive to the human mind and as successful at explaining variance as any known model, we may come up with new ways of conceptualizing our problems and the data we have collected on them.

## CONCLUSION

The purpose of this article is to introduce a computer program called NUANCE, the Naturalistic University of Alberta Nonlinear Correlation Explorer. We have contributed this program to the archival Web site of the Psychonomic Society ([www.psychonomic.org/archive](http://www.psychonomic.org/archive)), from where it may be freely downloaded. As its name suggests, NUANCE is a tool for exploring arbitrarily complex relationships between a set of predictor variables, and a dependent measure to be predicted. Written in Java in order to be platform-independent, NUANCE explores the space of possible relationships using genetic programming, which computes using natural selection. Our program makes use of a new age-weighted fitness function that is superior to existing published fitness functions. We have briefly introduced the concepts underlying genetic programming, explained the parameters and use of NUANCE itself, and presented some examples of its possible uses.

We believe that NUANCE is a useful tool for anyone faced with understanding complex relationships in a data set with one or more predictors and a single dependent variable.

## REFERENCES

- ANDREWS, S. (1997). The effect of orthographic similarity on lexical retrieval: Resolving neighborhood conflicts. *Psychonomic Bulletin & Review*, *4*, 439-461.
- FRAUMENI, J. F., JR. (1968). Cigarette smoking and cancers of the urinary tract: Geographic variation in the United States. *Journal of the National Cancer Institute*, *41*, 1205-1211.
- HOLLAND, J. (1992). *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. Cambridge, MA: MIT Press.
- KOZA, J. (1992). *Genetic programming: On the programming of computers by means of natural selection*. Cambridge, MA: MIT Press.
- KUSHCHU, I. (2002). Genetic programming and evolutionary generalization. *IEEE Transactions on Evolutionary Computation*, *6*, 431-442.
- LANGDON, W. B. (2000). Quadratic bloat in genetic programming. In D. Whitley, D. Goldberg, E. Cantu-Paz, L. Spector, I. Parmee, & H. Beyer (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)* (pp. 451-458). San Francisco: Morgan Kaufmann.
- LANGDON, W. B., & POLI, R. (1997). *Fitness causes bloat* (Tech. Rep. No. CSRP-97-09). Birmingham, U.K.: University of Birmingham, School of Computer Science.
- LUKE, S., & SPECTOR, L. (1998). A revised comparison of crossover and mutation in genetic programming. In J. Koza, W. Banzhaf, K. Chellapilla, K. Deb, M. Dorigo, D. B. Fogel, et al. (Eds.), *Genetic programming 1998: Proceedings of the Third Annual Genetic Programming Conference* (pp. 51-84). San Francisco: Morgan Kaufmann.
- MOORE, F. W., & GARCIA, O. N. (1997). A new methodology for reducing brittleness in genetic programming. In E. Pohl (Ed.), *Proceedings*

of the National Aerospace and Electronics 1997 Conference (pp. 757-763). Los Alamitos, CA: IEEE Press.

SOULE, T., & FOSTER, J. A. (1998). Removal bias: A new cause of code growth in tree based evolutionary programming. In D. Fogel (Ed.), *1998 IEEE International Conference on Evolutionary Computation* (pp. 781-786). Los Alamitos, CA: IEEE Computer Society.

SOULE, T., & FOSTER, J. A. (1999). Effects of code growth and parsimony pressure on populations in genetic programming. *Evolutionary Computation*, *6*, 293-309.

STONE, M. (1974). Cross-validation choice and assessment of statistical predictions. *Journal of the Royal Statistical Society: Series B*, *36*, 111-147.

STREETER, M. J. (2003). The root causes of code growth in genetic programming. In C. Ryan, T. Soule, M. Keijzer, E. Tsang, R. Poli, & E. Costa (Eds.), *Genetic programming: Proceedings of EuroGP 2003* (pp. 449-458). Berlin: Springer.

WESTBURY, C., BUCHANAN, L., SANDERSON, M., RHEMTULLA, M., & PHILLIPS, L. (2003). Using genetic programming to discover nonlinear variable interactions. *Behavior Research Methods, Instruments, & Computers*, *35*, 202-216.

WOLFRAM RESEARCH, INC. (2003). *Mathematica (Version 5.0)* [Computer program]. Champaign, IL: Author.

**ARCHIVED MATERIALS**

The following materials and links may be accessed through the Psychonomic Society's Norms, Stimuli, and Data archive, [www.psychonomic.org/archive](http://www.psychonomic.org/archive).

To access these files or links, search the Archive for this article using the journal (*Behavior Research Methods*), the first author's name (Hollis), and the publication year (2006).

FILE: Hollis-BRM-2006.zip

DESCRIPTION: The compressed archive file contains five files: dataset.txt, sample data set for modeling.

function.txt, list of functions enabled for the model.

GraphingTemplate.xls, example for graphing model output.

nuance.jar, java source for the model.

settings.txt, control settings for the model.

AUTHOR'S E-MAIL ADDRESS: [chrisw@ualberta.ca](mailto:chrisw@ualberta.ca).

AUTHOR'S WEB SITE: [www.ualberta.ca/~chrisw](http://www.ualberta.ca/~chrisw).

**APPENDIX**

**The Two Best Solutions, in Prefix and Standard Notation, as Evolved for a Problem (Considered in Example 1) Relating ON and Word Frequency to Lexical Decision RTs for Words**

Greedy Overselection best solution ( $|r| = .52$ ):

(/ (abs 2.0) (+ (+ (/ (abs (+ (- OFREQ (- (sqrt (sqrt (sqrt (abs 2.0)))))) (- (+ (+ OFREQ (abs (+ (sqrt (sqrt OFREQ)) 9.0))) (+ (+ (sqrt ON) OFREQ) ON)) (+ 0.0 (+ 9.0 OFREQ)))))) (sqrt OFREQ))) (+ (- OFREQ (- (+ (/ OFREQ (+ (+ (- OFREQ OFREQ) 9.0) 7.0)) (+ (+ (abs OFREQ) OFREQ) (sqrt 7.0))) (+ 0.0 (+ 9.0 9.0)))) 9.0)) (+ ON OFREQ)) (+ 9.0 OFREQ)), or

$$2 \text{ OFREQ} + \text{ON} + \frac{\text{Abs} \left[ -10.09 + \sqrt{9 + \sqrt{\text{OFREQ}}} + \sqrt{\text{OFREQ}} + 2 \text{ OFREQ} + \sqrt{\text{ON}} + \text{ON} \right]}{24.35 - 1.06 \text{ OFREQ}}$$

Median Pass selection best solution ( $|r| = .48$ ):

(/ 1 (+ ON (+ OFREQ (+ 10.0 OFREQ))))), or

$$\frac{1}{10 + 2 \text{ OFREQ} + \text{ON}}$$

Note: Both estimates correlate highly with each other and with the best estimate reported in Westbury et al. (2003), which was evolved using a different GP system. We have reproduced the raw NUANCE output here for illustrative purposes, as well as simplifications derived from Mathematica. The first equation was also simplified by hand. See the text for further discussion.