Design of Fuel Additives Using Neural Networks and Evolutionary Algorithms

Anantha Sundaram, Prasenjeet Ghosh, James M. Caruthers, and Venkat Venkatasubramanian
Laboratory for Intelligent Process Systems, School of Chemical Engineering, Purdue University,
West Lafayette, IN 47907

It is difficult and challenging to design high-performance fuel additives in an industrial-design setting where data are sparse and noisy, and fundamental knowledge is often limited. An automated framework is presented for the design of such fuel-additive molecules that minimize the intake-valve deposit in the automobile. A hybrid model that combined functional descriptors from a first-principles degradation model with a statistical/neural-network model was developed to predict additive performance, given the additive structure. The results of the predictive model are discussed for different real industrial case studies. An evolutionary method using specialized representation and constrained operators to enforce formulation constraints was used to generate optimal additive molecules that meet desired performance criteria. The evolutionary design strategy in combination with the hybrid prediction model was successful in identifying novel additive molecules that also possess good synthesis potential.

Introduction

A fuel additive is a substance added to gasoline in small quantities to provide improved performance or to correct deficiencies. Typically, gasoline additives often tend to provide these benefits at a lower cost and without the need for improvements in the refinery (Gibbs, 1990). As automobile designs have grown in sophistication over the years and have to meet several different performance criteria, the fuels also have to meet varied criteria for their performance in the automobile. Hence, additives that impart different functional characteristics to the fuel continue to be used. Gasoline additives historically have been used as combustion modifiers, antioxidants, corrosion inhibitors, anti-icing components, as well as deposit-control detergents (Gibbs, 1990; Kalghatgi, 1990). Among deposit controlling additives, the two primary ones are (1) intake-valve-deposit (IVD) controllers, and (2) combustion-chamber deposit controllers. In this article, we are specifically interested in the design of fuel additives that control the deposit formation on the intake valves of the automobile. Figure 1 shows the intake valve and its surrounding components in an automobile. The intake valve forms the opening into the combustion chamber. The fuel-injection nozzles spray gasoline directly on the intake valve. When the valve opens, it draws a mixture of fuel and air into the combustion chamber, where it is burned to supply power to the automobile. Over a period of time, both during and shortly after engine operation, deposits tend to form on the surface of the intake valve (Graham and Evans, 1992; Kalghatgi, 1990; Lacey et al., 1997), which affects driveability (Graham and Evans, 1992), cold start efficiency (Grant and Mason, 1992), power, acceleration, and knock characteristics (Arters et al., 1997), and emissions (Houser and Crosby, 1995). The severity of the implications on automobile operating characteristics due to deposit formation on the intake valve has necessitated the formation of the Coordinating Research Council (CRC) Intake Valve Deposit Committee in 1990 to address these specific issues. Additionally, the Environmental Protection Agency (EPA) has instituted a stringent test that every fuel package (gasoline + additives) now needs to clear before approving commercialization (EPA, 1996). This is the BMW-IVD test (ASTM, 1995), which involves driving a 4-cylinder 1985 BMW vehicle over the road for a total cumulative distance of 10,000 miles (16,000 km). The test cycle consists of 10% city, 20% suburban, and 70% highway mileage with an overall average speed of 45 mph (72 km/h). At the end of the test, the intake valves are removed and weighed. In order for
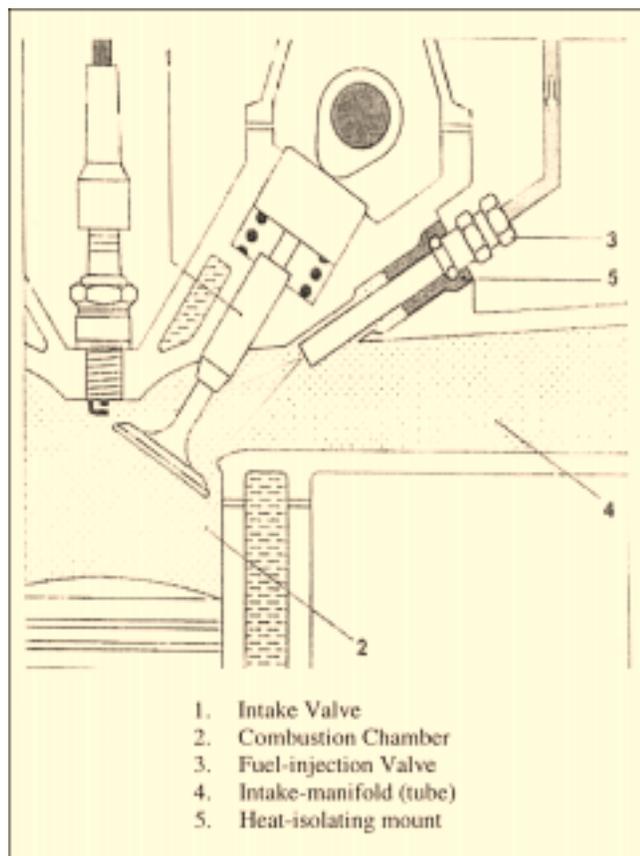
Figure 1. Intake-valve and manifold.

1. Intake Valve
2. Combustion Chamber
3. Fuel-injection Valve
4. Intake-manifold (tube)
5. Heat-isolating mount

"Given the characteristics of a particular grade of gasoline, what is the best structure of the fuel-additive molecule that will ensure that the overall fuel package (gasoline + additive) not produce intake-valve deposit more than a specified amount."

In order to address the design problem at hand, it is instructive to consider the mechanism of deposit formation on the intake valve. The mechanism has been hypothesized to be due to one or more of the following (Lacey et al., 1997; Bailey et al., 1995):

1. Autooxidation: Self-catalyzed oxidation during long-term storage.

2. Thermal Oxidation: Fuel-flow over hot surfaces leading to deposits and observed in operating equipment away from the combustion zone.

3. Pyrolysis: Decomposition of fuel and thermal-oxidative deposits on very hot surfaces, typically close to the combustion zone.

The actual governing reactions that lead to the formation of these depositions are quite complex, and a detailed description is not currently available. The key factors affecting the mechanism of IVD formation include the nature of the fuel and the operating conditions. In addition, other factors, such as valve and injector design (ChengShi Wai, 1992), automobile make (Homan, 1997), temperature gradients in the valve (Grant and Mason, 1992; Daneshgari et al., 1989), as well as engine-oil chemistry (Mitsui, 1993), have been found to have some effect on deposit formation. The quantities that are usually measured during engine tests are the operating conditions (valve temperatures, rpm, and so on), fuel grade, and in some cases, the chemical nature of the deposits. Hence, one usually does not have all the fundamental information required to develop a purely first-principles model that would predict the IVD formation, given the initial chemistry of the gasoline package (fuel + additives) and the operating conditions. Moreover, the engine tests are expensive and cost about $8000 a test. Hence, the data from engine tests are rather limited and, due to the nature of the tests themselves, quite

the fuel-package to successfully pass this test, it must produce an average deposit of less than 100 mg/valve (Lacey et al., 1997). This is the property of primary concern to the fuel-additive industry. Therefore the design problem is
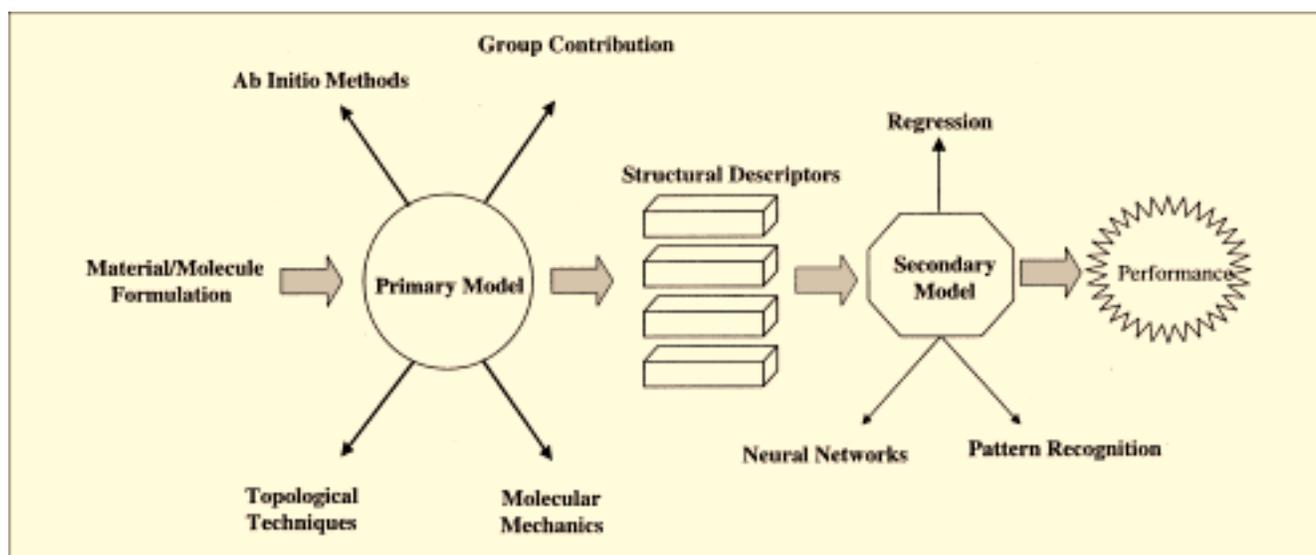
Figure 2. General framework of performance predictors from structure.
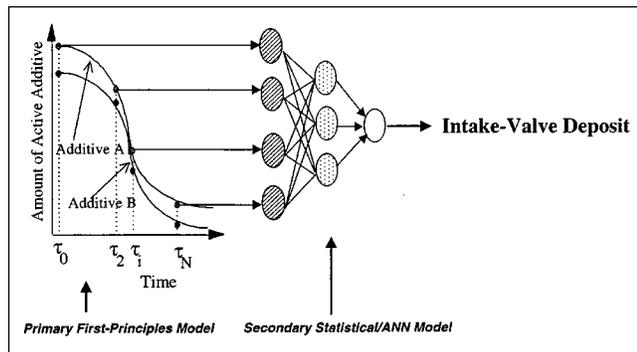
Figure 3. Hybrid neural network for IVD prediction.

Figure 4. Generic structure of fuel additives.

(a) Additive with single linker and tail; (b) additive with single linker but multiple tails; (c) additive with multiple linkers and multiple tails.

noisy. It is known that the car-to-car variation forms the dominant part of the repeatable error in these tests (Arters et al., 1997). These factors overrule a completely data-driven/statistical model for IVD prediction, since such a model will be prone to the noise in the data and might not capture, even qualitatively, the important phenomenological aspects of the underlying mechanism. Under such a scenario, a first-principles model in synergy with a data-driven/statistical model for performance prediction seems most promising. Figure 2 shows such a forward approach to property prediction from structure. Within this framework, if one were to predict performance instead of property, the structural descriptors would need to be augmented with phenomenological knowledge. This implies that *functional* descriptors that capture the underlying physics maximally need to be used instead of the so-called *structural* descriptors. With this mind, a hybrid (first-principles + statistical) model was developed to provide the right medium between computationally realizable physical relevance and predictive accuracy. The first-principles model captures in as much detail as possible the various phenomena that are hypothesized to be at play in determining the performance of the additive. This model determines, in the form of descriptors, the relative effects of the structural components of the additive on its performance. A statistical/neural network (NN) model then correlates these different descriptors to the measured IVD data. This architecture is outlined along the lines of the general forward model framework in Figure 3. The development of the first-principles model, which in this case is a solubility-based model for the fuel additive, and the use of this model to obtain input descriptors are described in the next section. These descriptions are followed by the development of the statistical/NN model that uses these functional descriptors as inputs and predicts the IVD. The development and validation of these models for two different kinds of engine tests, that is, the BMW test described earlier, and the Honda engine test, are also discussed.

First-Principles Model for Additive Stability

In the previous section the mechanisms underlying the formation of intake-valve deposits were outlined. In all these mechanisms there were two important features. The first concerns the stability of the deposit-forming precursors. Specifically, it is generally assumed that the deposits are not volatilized at the temperatures around the intake valve. Secondly, it is assumed that the insoluble fraction of the fuel falls out of solution unless stabilized by an additive, which itself may thermally degrade. It is hypothesized that these deposit-forming precursors can directly adhere to the intake-valve surface or undergo chain polymerization reactions with other species and form long-chain deposits on the surface. In either case, the thermal stability and the insolubility of the additive species play a crucial role in deposit formation. The primary function guiding the formulation of the fuel additive is its ability to scavenge the deposit-forming precursors from the high-boiling fractions of the fuel and prevent them from depositing on the intake valve. To impart this function it is imperative that the additive be thermally stable in the fuel milieu as long as possible. Thus, the first-principles model was focused toward determining the stability of the additive from its molecular structure and operating conditions.
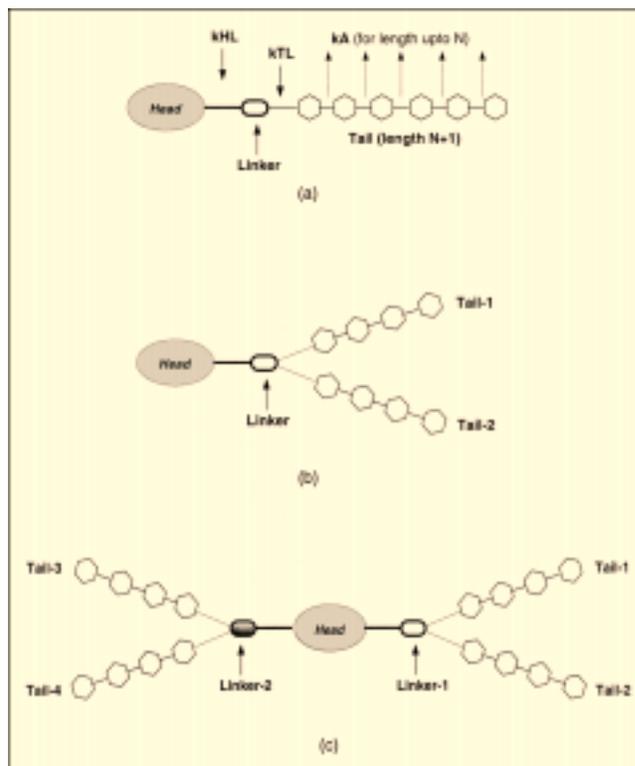
Figure 4 shows the functional components of a fuel-additive structure. The figure is a lumped functional look at the additive structure and is not a reflection of its detailed chemistry. The approach taken is hierarchical, where the functional aspects are broken down into the relevant structural aspects, and more details are added as necessary. The additive contains three main components:

1. *Head:* A set of (generally polar) functional groups collectively known as the head. The purported function of the head is to act as a receiver or scavenger of the high-boiling deposit precursors in the fuel. The head is generally the most

thermally stable component of the additive.

2. *Tail:* A long-chain component called the tail. The primary function of the tail is to keep the additive in solution long enough to be washed away by the fuel. In general, since the head is nonpolar and usually insoluble in the fuel, the tail helps the additive remain in solution, allowing the head to access the deposit precursors.

3. *Linker:* A transitional component that connects the polar head to the nonpolar soluble components, called the linker. The linker adds more stability to the additive and holds together the polar head with the soluble tail.

This structure of the additive can be described in several ways; however, the head−linker−tail description provides an accurate and fairly general representation of additives purely in terms of their key functional components. While in solution in the fuel + oil mixture, the different components of the additive may undergo various degradation reactions that alter their structural, and hence functional stability. While some components such as the head may remain stable and bounded to deposit precursors, the polymeric tail is relatively less so and tends to shorten in length, progressively decreasing the solubility of the additive in the fuel mixture. This degradation can reach a state where the additive molecule ceases to be effective as a detergent. It is clear that different mechanisms are at play here. The interaction between these different phenomena provides a dynamic solubility distribution for the additive in the surrounding fluid milieu. If the nature of this distribution and how it degrades from its initial state can be tracked, it may be possible to correlate to the IVD performance of the additive. The objective of the first-principles modeling is not so much to capture exactly in quantitative detail the different mechanisms involved, but to acknowledge the physics behind the relevant mechanisms to get a relative ordering of the performance of different additives. The main steps in the first-principles model development are the following:

1. Determine the tail-length distribution as a function of time, from the relative degradation rates for the tail, head, and linker.

2. Determine the solubility distribution (in terms of a standard measure) of the additive from its chain-length distribution.

3. Using information of the fuel characteristics, determine as a function of time, the fraction of *active* or solubilized additive.

The final step determines for each additive, a set of descriptors (so-called activity at different time lines) that can then be used as inputs to a neural network or statistical model to correlate the IVD performance of the additive. Figure 3 shows such a hybrid architecture.

Modeling Additive Degradation

The simplest generic structure of the fuel additive shown in Figure 4a is the basis of all calculations. The first step is to identify the potential sites of breakdown across the different components of the additive. The breakdown sites of interest in the functional components of the additive are the different positions in the polymeric tail, the tail-to-linker connection, and the head-to-linker connection. This functional description leads to species of the *head-only* type if the head-to-

linker bond is broken, the *head−linker−only* type if the linker-to-tail bond is broken, and the *head−linker−tails* type if the breakdown occurs at any of the different potential sites on the polymeric tail. It is obvious that there will be a distribution of species of the *head−linker−tails* type having different tail lengths. This allows us to provide a mathematical description that keeps track of all the temporal evolution of these different species that can result depending upon the site of bond breakdown. Without making any assumptions on the nature of the reactions or reaction mechanisms, the mathematical model governing the degradation of an additive structure given in Figure 4a, is as follows

$$\frac{\partial X_H}{\partial t} = f_H^{\text{form}}\left(X_{HL}, X_{j>0}, S_0 \right) - f_H^{\text{deg}}\left(X_H, S_0 \right)$$

$$\frac{\partial X_{HL}}{\partial t} = f_{HL}^{\text{form}}\left(X_{j>0}, S_1 \right) - f_{HL}^{\text{deg}}\left(X_{HL}, S_1 \right)$$

$$\frac{\partial X_i}{\partial t} = f_i^{\text{form}}\left(X_{j>i}, S_i \right) - f_i^{\text{deg}}\left(X_i, S_i \right), \ i = 1 \ldots N \quad (1)$$

In the preceding equation, X_H is the concentration (or alternately, the mol fraction) of the additive structures that contain only the head without any linkers or tails; f_H^{form} is the total rate of all reactions leading to the formation of such *head-only* structures. A head-only structure can be formed from all additive structures of different lengths if the head−linker bond is cleaved. In general, this rate could be a function of the concentration of the longer structures (head + linker, head + linker + tail of length 1, head + linker + tail of length 2, and so on) as well as the concentrations of other extraneous species (S_i) involved in these reactions. On the other hand, f_H^{deg} is the rate at which the head itself is degraded by reactions with other species S_0. The rate of change of concentration of the other additive species, such as X_{HL} (head + linker) and $X_{i>0}$ (head + linker + tails of length 1 and greater), can be modeled in a similar fashion. Since the nature of the particular reactions can be very complex, it is difficult to model these mechanisms in great detail. Since an exact kinetic model was not sought and the ultimate aim was to model the relative stability of the additives, the following assumptions were introduced:

1. The degradation reactions were modeled as first order and irreversible.

2. The rate of degradation was parameterized by three rate constants: (a) rate constant, k_{HL}, between the head and linker; (b) rate constant, k_{LT}, between the tail and the linker; and (c) rate constant, k_T, along the length of the tail.

3. The tail-length distribution [for molecules of the form head, head + linker, head + linker + tail (lengths 1 to N), etc.] was initially assumed to be normal with a given mean and standard deviation. These were in turn obtained from the average molecular weight and dispersity data of the initial distribution.

Even if the model simplified the reaction mechanisms, it will at least effectively order the thermal stability of various additives, which is all that is needed for a functional description. Under the preceding assumptions, Eq. 1 can now be rewritten as a series of linear ordinary differential equations

(ODEs).

$$\frac{dX_H}{dt} = k_{HL}\left(X_{HL} + \sum_{i=1}^{N} X_i \right)$$

$$\frac{dX_{HL}}{dt} = -k_{HL}X_{HL} + k_{LT}\sum_{i=1}^{N} X_i$$

$$\frac{dX_1}{dt} = -(k_{HL} + k_{LT})X_1 + k_T\sum_{i=2}^{N} X_i$$

$$\frac{dX_j}{dt} = -[k_{HL} + k_{LT} + (j-1)k_T]X_j + k_T$$

$$\sum_{\substack{i=j+1 \\ \forall\, 1 \le j \le N-1}}^{N} X_i$$

$$\frac{dX_N}{dt} = -[k_{HL} + k_{LT} + (N-1)k_T]X_N, \qquad (2)$$

where N is the maximum tail length in the initial distribution, t is the time, and X_i is the concentration/mol fraction of a tail of length i (where i varies from 1 through N). In order to solve Eq. 2, one now needs to specify the initial conditions. When the tail is monodisperse, the initial condition is specified in a straightforward fashion, setting the longest tail concentration to one and concentration of all shorter chains to zero. When the tail is polydisperse, the initial conditions for the preceding system can be generated from the information of the polydispersity index and average molecular-weight data. We have assumed this distribution to be normal. The solution of the preceding equation will yield the mol fractions of all the different species mentioned previously, as a function of time. The preceding equations were made dimensionless in time by setting $\tau = k_T t$. This leads to two parameters that control the degradation behavior of the system, namely k_{HL}/k_T and k_{LT}/k_T. However, the actual reaction mechanisms are rather complex, and hence it is difficult, if not impossible, to determine the rate constants. Thus, an approximation to the ratios was obtained by the following procedure:

1. The most reactive bonds in the structure were identified at the appropriate positions on the tail, between the tail and the linker, and between the linker and the head. This was based on the relative reactivity of different functional groups and their stability. The bond dissociation energies of these bonds would be representative of the activation energy for the degradation of different components of the additive.

2. Small molecule analogs were constructed for each of the bonds whose dissociation energies had to be estimated. These analogs were constructed in such a fashion as to retain a similar charge behavior around the bond under consideration. Functional groups and atoms farther along the chain from the bond of interest were discarded and the molecule was truncated with a hydrogen or methyl group at an appropriate position to obtain the analogs.

3. For the analog, the bond dissociation energies of the bond considered were calculated by using semiempirical quantum chemical calculations. This involves estimating the heats of formation for the molecule as well as the two radicals formed after the bond was broken. The bond dissocia-

tion energy is the difference between the sum of the heats of formation of the radicals and the heat of formation of the analog.

4. The ratio of the rate constants for an irreversible, first-order degradative reaction was calculated as the ratio of the exponentials of the calculated bond dissociation energies (from the Arrehenius rate equation)

$$\frac{k_{HL}}{k_T} = \exp-\left(\frac{(E_{HL} - E_T)}{RT} \right), \qquad (3)$$

where E_{HL} and E_T are the bond dissociation energies of the appropriate small molecular analogs of the linker–head bond and the polymeric tail bond, respectively; T is the temperature of evaluation; and R is the universal gas constant. The operating temperature for this study was 523 K (250°C). This is the typical temperature of the surface of the intake. With the rate constants now specified, a vectorized MATLAB code was developed to solve the system of equations in Eq. 2.

Multiple Tails

Most of the structures of actual interest include heads with multiple linkers and tails, as shown in Figures 4b and 4c. A single head portion is attached to multiple tails through a single linker in Figure 4b and by different linkers in 4c. In order to simulate the tail-length distribution of the structure in Figure 4b, the tail length is defined to be the sum of the lengths of all tails on the linker. For multiple tails greater than two, the approach is to first perform the simulation for two tails and then successively add in the distribution of the other tails. The simulation approach presented for the single-tail case is extended to the two-tail scenario by using the following ideas:

1. The degradation of any single tail is independent of all other tails attached to the same linker + head.

2. The total concentration of heads (attached to linker and tails, as well as separate) is a conserved quantity irrespective of the number of linkers + tails present, since the rate of formation of single heads depends only on the rate of degradation of the head–linker bond(s).

Under these assumptions, the concentration of a structure with a total "linker-A + tail-A of length i" with a second "linker-B + tail-B of length j" is given by Eq. 4.

$$X_{ij} = (1.0 - X_H)P_i P_j \qquad (4)$$

In this equation, X_{ij} is the concentration of the structure just mentioned; X_H is the concentration of structures with heads only; and P_i and P_j are the probabilities of finding single tails of length i, and length j, respectively. They are determined independently by performing two simulations with single-tail generic structures of the kind shown in Figure 4a. Equation 4 follows from the assumption of the independent degradation of each individual chain attached to the head. Once the preceding concentration is evaluated for each combination (i, j), the distribution of the sum (given by X_s, where $s = i + j$) is given as

$$X_s = \sum_{i+j=s} X_{ij}. \qquad (5)$$
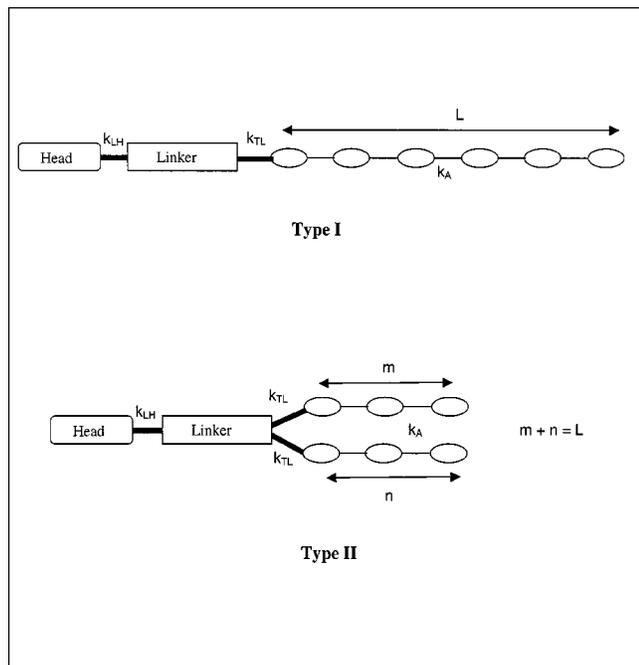
Figure 5. Type I and Type II structures.

Figure 7. Comparison of chain-length distributions for Type I and Type II structures.

From a computational point of view, this is easily achieved by performing two complete simulations (when both tails are different) to obtain, $X_1(t)$ and $X_2(t)$. They correspond to the tail-length distribution for each of the two tails. At each instant of time, the outer product of the two vectors yields the matrix of the joint distribution of the two tails. Summing up the contradiagonals of the resulting matrix gives the concentration of various possible sums of the two tails.

In order to obtain a qualitative verification of the tail distribution behavior for typical cases, simulations were performed to compare a two-tailed additive and its single-tail counterpart that would possess the same total length. The

two additives, whose chain-length distributions were simulated, were structurally similar except for the number and manner in which the tails are connected to the head. The structures considered are shown in Figure 5. Type I is a single-tailed structure, and Type II is a two-tailed one. For this particular case, the rate constants governing the degradation of the head from the linker and the linker from the tail were very low for both the additives. This implies that the head–linker and tail–linker bonds were very strong, and tails were never lost completely by the breaking of these bonds. The only difference between the additives was the presence of an additional tail in the two-tailed structure and its connectivity. Figure 6 shows the concentration vs. time plots for three different values of the chain length (l) for Type-I (single-tailed) and Type-II (two-tailed) structures. In both cases the concentration of these structures shows a steady increase because the head–linker and tail–linker bonds are always very stable, that is, the first tail unit is never lost. Hence the longer tails are successively degraded to form shorter and shorter tails. The difference between Type I and Type II structures for the zero tail-length ($l = 0$) case is due to the fact that Type II has longer tails in general but a lower initial

Figure 6. Comparison of degradation behavior of different chain-length additives.
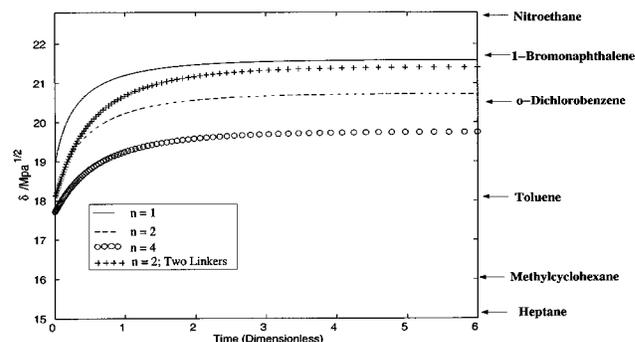
Figure 8. Solubility parameter vs. time for additives with different structures.

concentration of very short and very long tails. This means that the tail degradation to structures with $l = 0$ is much slower in Type II. This difference in the transient behaviors of the tail distribution in Type I and Type II is shown in Figure 7. It is evident from the plots that two-tailed (Type II) structures tend to keep the tails longer and retain a higher concentration of species with tails compared to the Type I structures. This leads to the conclusion that the alkane-like solubility characteristics imparted by the tail would be sustained for a longer duration in the case of the additive with two short tails, as compared to the additive with a single long tail.

Solubility Descriptor Determination

At this stage, we have a degradation–kinetics model that predicts the temporal degradation of the additive structures. This is a necessary first step, but what we are truly interested in is a phenomenological descriptor that could be correlated to IVD. In our problem, this is the *solubility descriptor*. Typically, solubility is indicated by the cohesive energy density, which in turn is characterized by the Hildebrand parameter (Hildebrand and Scott, 1962). This is a measure of the internal-energy density and represents the amount of energy required to move two molecules of a species to infinite separation in solution. The Hildebrand parameter can be estimated by group contribution methods, of which the modified

Hansen's methods (Barton, 1991) were found to be the most suitable for the purposes of descriptor calculation in this case. These were obtained from the *Handbook of Solubility and Cohesion Parameters* (Barton, 1991). Hansen's group contribution method determines three separate contributions to the Hildebrand parameter, which include dispersion, polarity, and hydrogen bonding. A molecule is first split into a group of functional groups that have fixed contributions to each of these terms. For a molecule containing N_f functional groups, the Hildebrand parameter was estimated as follows (Barton, 1991)

$$\delta_d = \sum_{i=1}^{i=N_f} \frac{F_i^{zd}}{V_i}$$

$$\delta_p = \sum_{i=1}^{i=N_f} \frac{F_i^{zp}}{V_i}$$

$$\delta_h^2 = \sum_{i=1}^{i=N_f} \frac{\left(U_i^{zh}\right)^2}{V_i}$$

$$\delta = \sqrt{\delta_d^2 + \delta_p^2 + \delta_h^2}, \tag{6}$$

where δ_d, δ_p, and δ_h are the dispersion, polar, and hydrogen bonding contributions for the entire molecule; F_i^{zd}, F_i^{zp}, and
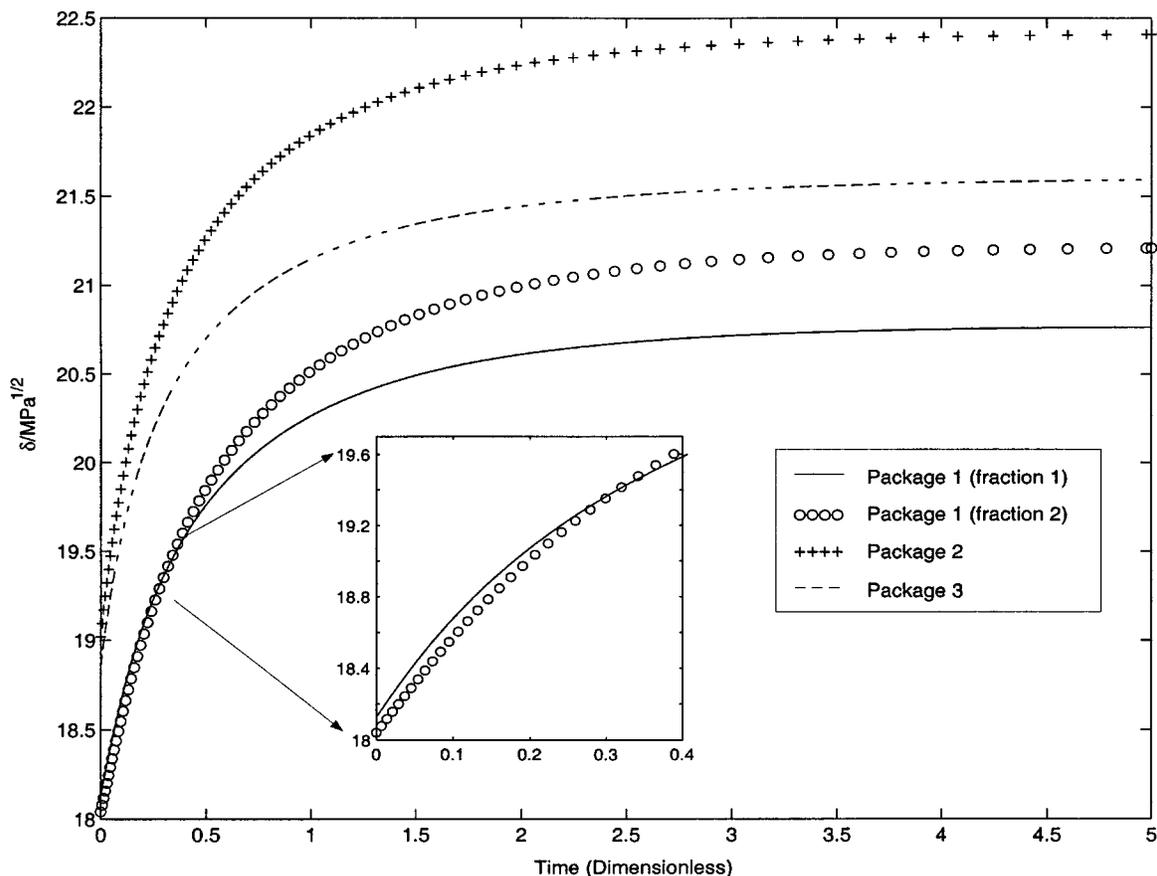
Figure 9. Solubility parameter variation with time for different additives (BMW database).

U_i^{zh} are the functional group contributions to each of those three terms, respectively; and V_i is the group-contribution to the molar-volume of the molecule, from functional group i. When a mixture/solution of two species (like that of a fuel−additive mixture) is considered, a small Hildebrand parameter of one species (say, A) compared to the other (say, B) indicates that A will not be able to provide enough energy to disperse B, and hence they will remain immiscible. Hence, a large difference between the Hildebrand parameters of the components indicates poor solubility of A in B. Fuel additives contain a distribution of chain lengths, and hence a distribution of Hildebrand parameter values. The Hildebrand parameter for the overall fuel additive was calculated as the mol-fraction weighted sum of the different fractions constituting the additive. Specifically

$$\delta_T = \sum_{i=1}^{i=N} X_i(t)\,\delta_i, \qquad (7)$$

where δ_i were the Hildebrand parameters estimated for each species, δ_T was the total Hildebrand parameter value, and X_i were the mol fractions of species of chain length, i. The

mol fractions were in turn distributions (varying with chain length) changing over time due to the degradative reactions.

Figure 8 shows the temporal (dimensionless) evolution of the Hildebrand parameter (δ) for different structures that vary in the number of tails attached to the linker, as well as a case where the head contains two linkers and two tails attached to each linker. These structures are based on an additive in the BMW database, which is a collection of 92 engine tests on different additives. Realistically, only one or two tails per linker are probable in the actual additive structures. For cases with greater than two tails, it was assumed that there was some way more tails could be attached to the linker independently, without changing the kinetic stability characteristics of the linker. Figure 8 compares the Hildebrand parameter for different additive structures against the estimates of the Hildebrand parameter for alkanes, aromatics, and amines on the background. For example, nitroethane has a Hildebrand parameter value of 22.8, while toluene is 18.2. The plot indicates that as more tails break down to shorter lengths, the Hildebrand parameter tends to increase and settle down in the aromatic/naphthalene end of the spectrum. Also, as the number of tails increases, the additives tend to reside at lower values of the solubility parameter. This is because the
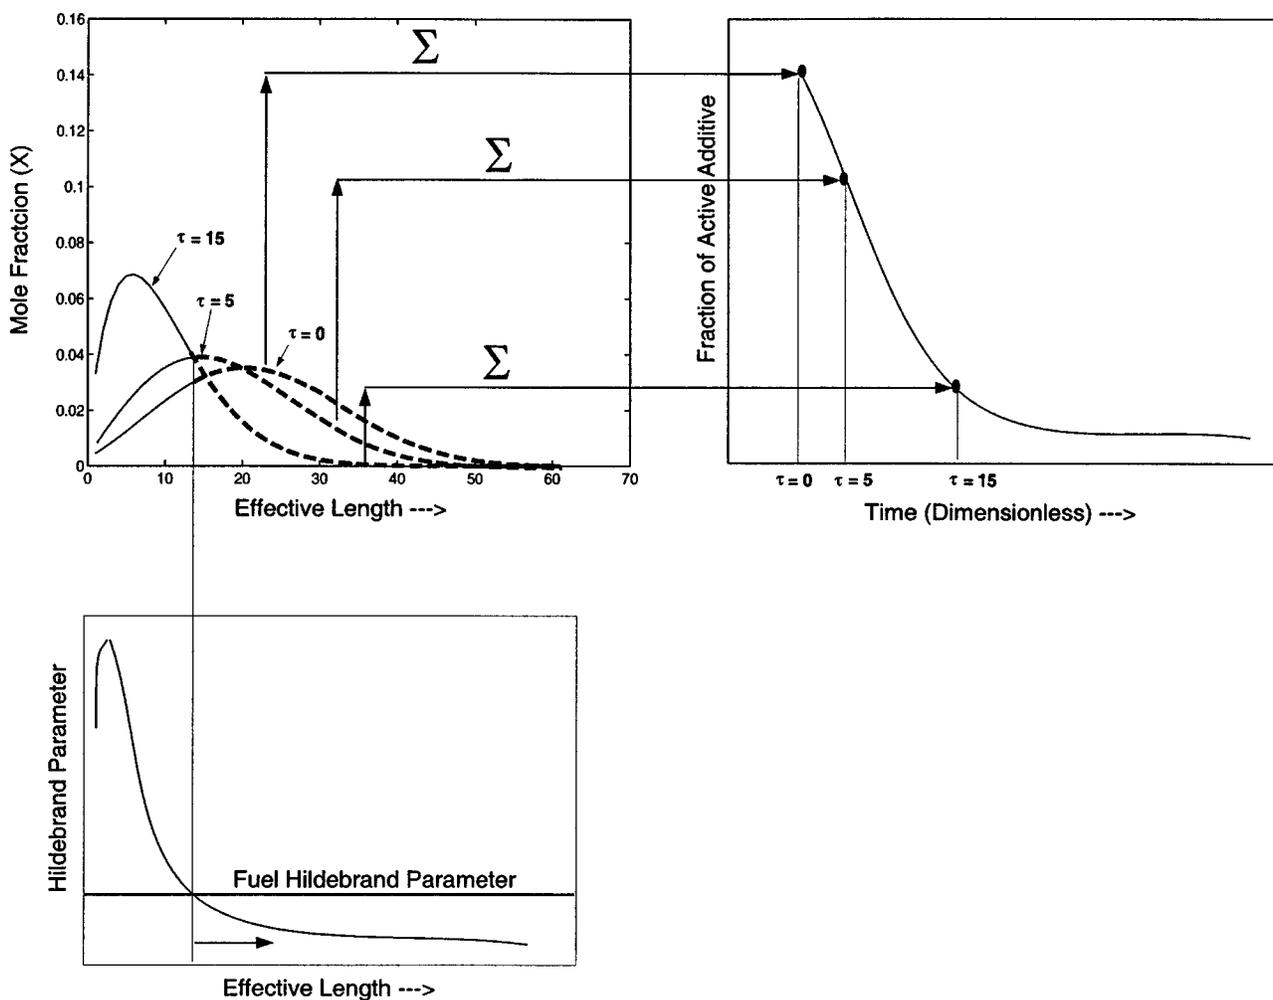
Figure 10. Determination of amount of active additive from solubility distribution.

head−linker and tail−linker bonds do not degrade thermally. Consequently, even after a long period of time the first tail unit remains in the structure. This indicates that one would get a greater nonpolar contribution from the first tail unit when more tails are present, thus lowering the final solubility parameter value. However, in the final case shown in the figure (four tails attached to the head via two linkers), there is an increase in the final solubility value compared to the case with four tails. This is due to the additional polar contribution of the extra linker. As long as the additive cohesive energy density was comparable to that of the liquid fuel, then the additive would remain soluble in the liquid portion of the fuel throughout the period under consideration. The polar head that binds to the deposit precursors would be more effective, as it was kept in solution longer.

To determine whether the solubility calculations lead to differences, and hence ordering between additive packages, four different additive packages from the BMW database provided by Lubrizol were tested. Due to the proprietary nature of the data, the actual structures in the additive package cannot be disclosed here. The main difference between fraction-1 and fraction-2 for additive package-1 was the presence of two additional tails connected to the head through an additional linker in fraction-2 of the package. As evident from the inset in Figure 9, initially fraction-2 has a slightly lower Hildebrand parameter value due to the presence of longer tails. The polar contribution of two linkers in fraction-2, as opposed to one in fraction-1, dominated the solubility contribution of the structure as the tail started degrading, and it finally settled at a much higher value compared to that of package 1. On the other hand, package 2 showed higher solubility parameter values. This was due to the presence of only a single tail and a large polar contribution of the head for this particular additive structure. Package 3 had a lower final

value due to a smaller polar contribution from the head compared to that of package 2. The solubility simulations led to the following conclusions:

1. The degradation of the tail played an important part in the dynamic behavior of the solubility characteristics.

2. The topology or the internal connectivity in the additive structure also played a crucial role in determining its degradative characteristics. Hence, they also dictated the nature of variation of the solubility characteristics with time.

3. The relative polarity of the head/linkers of different additives determines the magnitude of the final difference in the Hildebrand parameter values. This was mainly due to the greater thermal stability of the head and the linker.

Consider now a solvent milieu for the additive that possesses a predominantly alkane-like character (say, octane). This would be, for instance, largely characteristic of the initial composition of the fuel. The experimental estimate of Hildebrand parameter values for octane is about 15.4 $MPa^{-1/2}$ (Barton, 1991). In order for the additive to be soluble in the fuel, the difference between the Hildebrand parameters for the additive and the fuel should be within a specified limit. This difference is typically about 5 $MPa^{-1/2}$, that is, $|\delta_i - \delta_j| < 5$.

We discussed previously the estimation of the additive Hildebrand parameter from the additive structure and operating temperature (Eqs. 6 and 7). In addition, if the fuel Hildebrand parameter could be estimated, then the fraction of the additive distribution miscible in the fuel solvent could be calculated. Figure 10 shows the essential steps of such a calculation. First, the Hildebrand parameter of the additive as a function of the chain length was obtained using the group-contribution method described previously. This is not a function of time, and is shown for a typical case at the bottom of Figure 10. On this plot the fuel Hildebrand param-
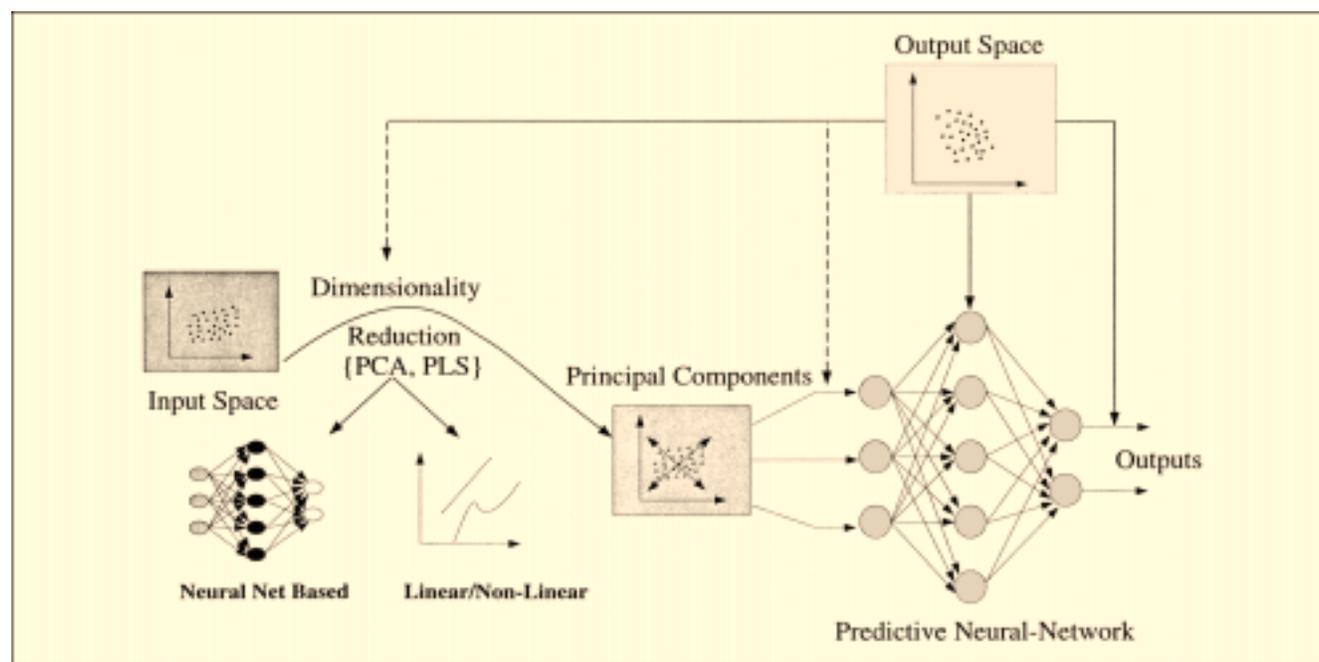
Figure 11. General architecture of neural networks considered in the BMW database.

eter is demarcated on the *Y*-axis. It was assumed that there are enough fuel components that are compatible with any solute/additive species with Hildebrand parameter values below this critical value. These species are indicated inside the square block on this plot. The top lefthand portion of Figure 10 shows the additive distribution at a given time *t*. From this plot, and using the information of the additive species/lengths that are compatible with the fuel, the distribution of the active or solubilized additive species was obtained. The sum of the molar concentrations of all these species gives one point on the activity plot, as shown in the righthand top portion of Figure 10. The dynamic activity curve can be constructed in this fashion, for a given additive package + fuel combination. This forms our *functional descriptor*, which is our input to the secondary model.

IVD Performance Predictions for Fuel Additives

The final and an important step in the performance prediction of intake-valve deposits is the determination of a regression model that correlates the functional descriptors to the IVD data from engine tests. Several different models can be developed at this stage. However, any model contains two main components, namely, the inputs to be used by the model and the model's internal parameters. The inputs in all cases were the amounts of active additive at different times. These times were not fixed *a priori*, but optimized toward best prediction. This was done in a sequential but trial-and-error fashion. The influence of the actual inputs chosen on the prediction accuracy of the model is less than that of the model's internal parameters. While a wide variety of statistical models may be explored for this purpose, we restricted ourselves to linear models based on projections to latent structures (PLS) (Geladi and Kowalski, 1986) and neural networks (Haykin, 1999). The main reason for the examination of linear model is their simplicity. On the other hand, neural networks provide general nonlinear architectures for function approximation and pattern recognition. The complexity of the underlying phenomena and the lack of accurate measurements make the quantitative information built into the first-principles model highly difficult, if not impossible. This and the noise in the engine test data make neural networks ideal candidate models for IVD prediction. Figure 11 shows the general architecture of neural-network models examined for the first case-study involving the BMW engine-test data. While there are no restrictions on the direction of connectivity of neural-networks for a general case, attention here is focused only on feed-forward networks where connections go from inputs to outputs across different layers. These networks are widely used for prediction and pattern classification (Haykin, 1999).

The main components of the neural-network architecture are the inputs, the number of hidden neurons, and the nature of the transfer functions. Developing a neural-network prediction model involves two phases: training and testing. During the training phase, the weights on the connections between the nodes are regressed from a given set of inputs and their corresponding output data. Several algorithms exist for training the weights of the neural-networks, such as back propagation (Rummelhart, 1986), conjugate-gradients (Charalambous, 1992), and general nonlinear optimization meth-

ods, such as the Levenberg-Marquadt (LM) technique. The LM method was used for training due to its added robustness to locally optimal weights (Hagan and Menhaj, 1994). However, it is expected that the use of conjugate gradients or any such technique for this purpose would not drastically alter the results. During the testing phase of the neural-network, the network with the trained weights is presented with data that were excluded from the training set. The performance of the neural network on data not seen before is indicative of its generalization or extrapolation performance.

The engine-test data that were available to train the neural networks were limited (a total of 100 points), given the expense of the tests. To reliably train models to predict data outside the training set, the cross-validation method was used. In all the cases, a single hidden-layer neural network was used, but the number of neurons was varied. The entire data set was partitioned into a training set and a testing set (approximately 10% of the total data). The neural network was trained on the training data and then tested to evaluate its performance on the test data to determine its root-mean-squared error (rmse) during the testing phase. Next, the test set was added into the data set and a different partition made, to obtain different training and testing data sets. The architecture was again trained and tested on this new partition. This procedure was repeated several times until each data point had appeared at least once in both the testing and training sets. The performance of the architecture is reported as the average rmse on this cross-validation (during the test phase).

BMW database

The BMW database consisted of 92 engine-test results after screening for outliers. The additive structures and fuel characteristics were provided for each engine test. The fuel Hildebrand parameter required for estimating the amount of active additive was not directly available. For the BMW database, a fuel profile was obtained based on the composition of the fuel and its performance in engine tests without any additives. All the fuels used in the engine tests were ranked based on this profile. A base value (which was later optimized) of the Hildebrand parameter was assigned to the most "severe" fuel that had the worst IVD performance (without additives). All the other fuels were assigned values relative to the most severe one on a linear scale with an adjustable slope. The slope was also optimized to give maximum correlation to the IVD performance. The development of the first-principles-based hybrid model from the engine-test data followed the steps given below:

1. The solubility and molecular weight distribution were determined as a function of time, as outlined in the first-principles model development.

2. The fraction of the additive package that remained active in the given fuel was determined. Activity was defined as the number of moles of the additive that had a Hildebrand value lesser than that of the fuel at any time.

3. The amount of active additive at different times gave the solubility descriptors that described the time-varying solubility characteristic of the additive in the given solvent.

For the BMW database, these descriptors were used as inputs to different neural-network architectures and the best

set was reported. The inputs were mean-centered and variance-scaled before feeding into the network. Several different models were examined, including some linear ones. The linear models were clearly outperformed by the neural-network architectures. Comparisons were made between networks using solubility descriptors as inputs and the best networks derived using structural descriptors as inputs. These structural descriptors were picked to reflect the structure of the additives in terms of various functional groups and were provided by the chemists at the Lubrizol Corporation. Although these descriptors did capture some of the important structural characteristics of the additive, the connection between these descriptors and the macroscopic property of interest, in this case, the amount of intake-valve deposit, was not transparent. On the contrary, the functional descriptor derived from first principles in the previous sections captures the inherent chemistry in a clear and transparent manner. The results are encouraging using this functional descriptor and are summarized in Figure 12. The best network in terms of accuracy (rmse on 10-fold cross-validation) and reliability (standard deviation of the error across different data partitions) was a radial-basis neural network that used a *single* solubility-based descriptor (the moles of active additive at time = 0). The performance of this network was comparable to that of the best network that used as inputs linear projections of 36 structural descriptors extracted directly from the structure of the additive, without the aid of any phenomenological model. The first-principles hybrid neural-network model overhelmingly outperformed the other models, especially considering that only a single-descriptor was eventually used. Primarily, this was due to the rich phenomenological

No. of Input Descriptors	Projections (PLS/PCA/ None)	NN Architecture Transfer Function [No of Hidden Neurons]	RMSE (mg) in Testing (CV)
36	None	Tan-Sigmoid [7]	273*
4	PLS	Log-Sigmoid [3]	105*
3[†]	None	Tan-Sigmoid [3]	142
7	PCA	Tan-Sigmoid [3]	200
6[††]	None	Tan-Sigmoid [3]	172
1**	None	Radial Basis [4]	124

*Large Standard Deviation; [†]Descriptors extracted from PLS factors
[††]Descriptors extracted from PCA factors
**Solubility Descriptor extracted at time $\tau = 0$

Figure 12. Comparison of models for IVD prediction: BMW database.

knowledge included in the input descriptor through the first-principles model as opposed to picking standard structural descriptors that one can obtain using any conventional molecular-modeling software.

Honda database

The Honda database represented the engine-test results of IVD performance based on a Honda two-valve generator. The Honda generator test is a shorter test compared to the BMW test, taking about 80 hours, and is generally considered to be
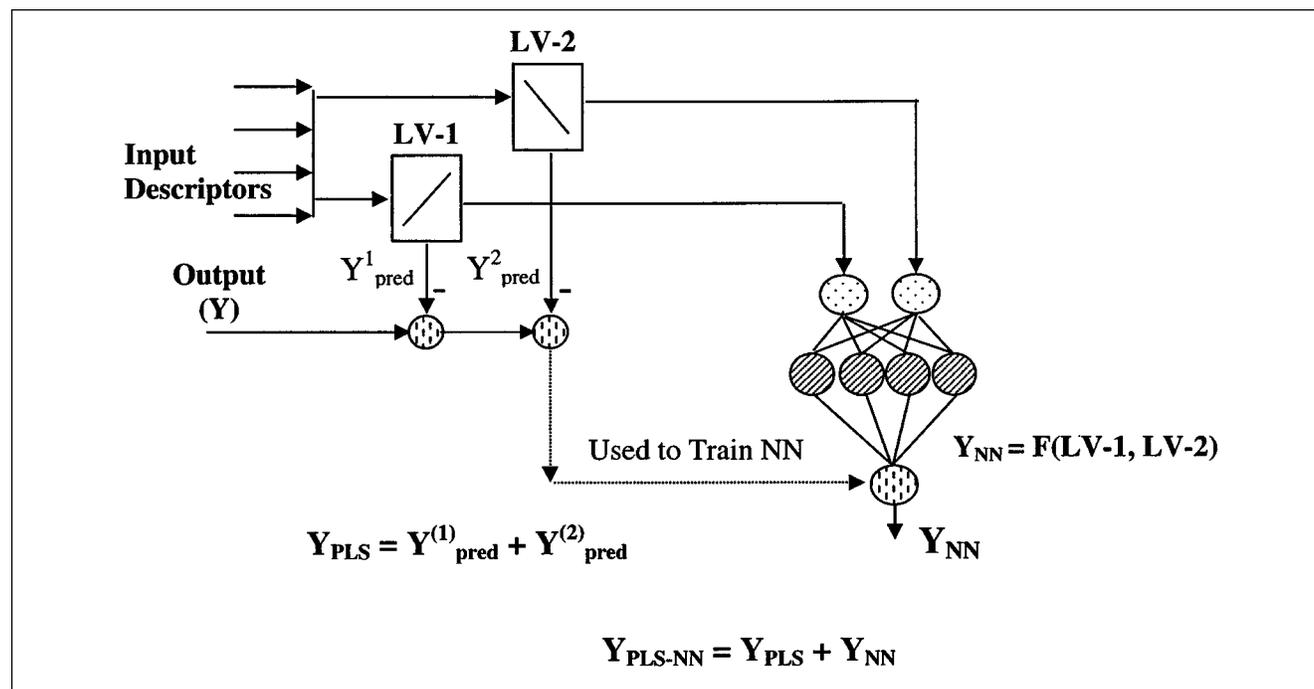
Figure 13. General architecture of a PLS-NN model.

more consistent statistically. Although, most of the additive components that formed the basis of this database were similar to those used in the BMW engine tests, there were some differences, the most important being the lack of information about the fuel Hildebrand parameter required for determining the amount of active additive in the fuel. As a reasonable estimate, therefore, the fuel Hildebrand parameter was allowed to vary between 18 and 28 MPa$^{-1/2}$ for the Honda database. Several time vs. active-amount curves were constructed for each additive, corresponding to values of the Hildebrand parameter given earlier. These curves were then used as inputs to a neural network. Since this leads to overparameterization, projection methods especially PLS-based techniques (Geladi and Kowalski, 1986; Qin and McAvoy, 1992) were used in series with the neural network for prediction.

Partial least squares (PLS) projection to latent structures are based on linearly projecting (or rotating) the input as well as the output space so that they are optimally correlated (in a linear sense). The projection part of PLS is called the *outer relationship*, and the correlation or prediction part of PLS is called the *inner relationship*. For the IVD prediction problem there was only one output variable, that is, the IVD data from engine tests. There was then no need for output projections. Originally PLS methods implied only the use of linear models for both inner and outer relationships. Over the years, nonlinear PLS methods have been developed using both nonlinear outer relationships, as well as nonlinear inner relationships. One approach in this regard is the generalized nonlinear PLS regression based on neural-networks proposed by Qin and McAvoy (1992). In this approach several neural networks were used in series to successively correlate each input factor with the corresponding output factor. Residuals from preceding neural-network relationships were used for successive training. A variant of this approach is to use the neural network only as a transformation relation for the output, that is, as the output outer relationship (Andersson et al., 1996). While a series of neural networks would clearly overparameterize the model for the given data, using the neural network only as a transformation of the output also may not be sufficient, considering there is only one dependent variable. In addition, for the latter method, a smooth nonlinear function in the residuals needs to be established. Given the noise in the data, this requirement is difficult to meet and discern. A better approach to the same problem would be to model the relationship between the input and output to consist of a linear and a nonlinear part. The linear relationship would be determined using the traditional PLS technique and the evaluated output residuals. The input factors, determined as a result of the first step, would then be correlated to the residuals via a neural network. This significantly reduces the number of parameters in the system, and using the neural network to "fit" only the nonlinear interactions in the data (if any) after the significant linear effects have been accounted for through the PLS model. This approach is outlined in Figure 13. Fitting the noise instead of the data trend is a possible source of error in this method. However, the networks were cross-validated several times to eliminate this possibility (Schenker and Agarwal, 1996). The architecture of this PLS-NN model would then consist of the number of PLS factors used, plus the architecture of the neural network itself. These

Best Linear Model: ($\tau = 0,1,5$; $\delta = 18\text{-}28$), 3 Latent Variables; RMSE: 33 mg in Testing (CV)

Best NN Model: ($\tau = 0,1,5$; $\delta = 18\text{-}28$), 2 Tan-Sigmoid Hidden Neurons RMSE: 35 mg in Testing (CV)

PLS-NN Models : RMSE in Testing Phase of CV					
Number of Hidden Neurons	Latent Variables				
	2	3	4	5	6
2	32.8	31.98	30.7	32.1	31.5
3	31.7	31.22	31.62	33.4	31.3
4	33.9	31.4	31.7	31.6	33.6

Figure 14. Performance comparison of different models (Honda database).

factors were determined by cross validation, as described for the BMW database. The performance of different PLS-NN models using the Honda engine-test data is shown in Figure 14. As the results indicate, the generalization performance is very consistent across different architectures. It is clear from the results that the model was able to perform very close to the quality of the given data. The Honda engine-test data had a smaller standard deviation and lower test-to-test variability. This was reflected in the fact that the rmse in cross validation across different models was consistently close to about 32 mg for the Honda database. The best PLS-NN model had an rmse of 30 mg in cross-validation. The experimental error between repeats for this test was around 25 mg.

Extension to Blends of Fuel Additives

The preceding model was extended to blends, which are fuel-additive packages that contain more than one fuel-additive molecule/component in them. As with the previous data, no information about the fuels was used, and the latent variables used in the models were projections of the descriptors calculated at different times while varying the fuel Hildebrand parameter between 18 and 28 MPa$^{1/2}$. Cross-validation was performed over 50 partitions, and the average rmse is reported in the table. It was evident that the accuracy of the best models did not vary significantly, with linear models slightly better than the best neural-network models. This may be due to the sparseness as well as the noise in the data. While the accuracy of the forward models in explaining the experimental data for the Ford model was not great, the solubility descriptors did capture quite significantly the formulation knowledge. Figure 15 shows the plot of the best solubility descriptor on the X-axis with expert assigned values for the thermal stability of the blends used in the case study. The two descriptions are closely correlated with a correlation coefficient of 0.965. This showed that, because of the nature of the hybrid-model, a clear separation between the underlying physics of the mechanism and the data-based statistics is achieved. This implied that even when the hybrid model failed to provide an accurate description of the IVD behavior of blends (due to the nature of the training data or other un-
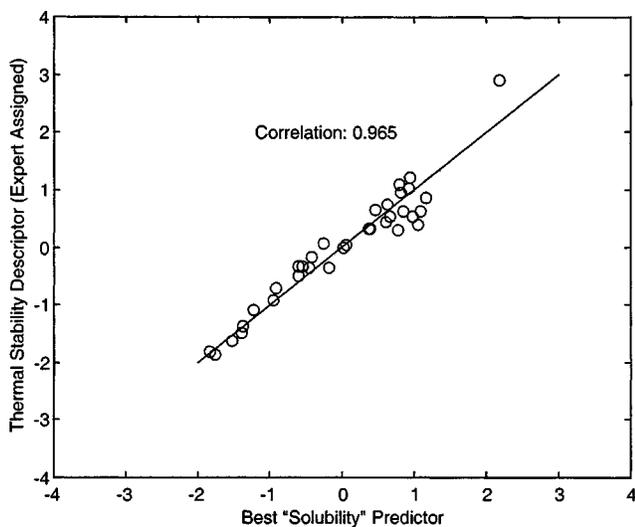
Figure 15. Correlation between expert assigned and solubility descriptors.

modeled effects), the first-principles model remained quite accurate in capturing the underlying formulation knowledge behind the stability of the additives. This is crucial, because even in the absence of engine test data, a screening procedure for additives based on stability could be established.

Fuel Additives Design Using Evolutionary Algorithms

Up to this stage, we have demonstrated that a hybrid approach utilizing fundamental models in conjunction with some statistical/neural-network algorithms provides improved predictions for a number of different case studies for a difficult design problem, such as fuel-additive design over conventional structure-property methodologies employing a multitude of ad hoc structural descriptors. This sets the stage for the inverse or reverse-design problem, which involves construction of fuel-additive packages that meet a desired IVD performance constraint. Although several techniques could be potentially applied towards the inverse problem, most of them rely strongly on the functional nature of the predictor method. In our particular case, the presence of the neural network in the hybrid model precludes the use of mathematical-programming methods, because the accompanying solution strategies do not work well with neural-network objective functions. Evolutionary methods offer a powerful alternative, due to their robustness to discontinuity and nonlinearity in the objective function. Therefore, it was the ideal method for the design of fuel additives using hybrid forward model for performance prediction. Genetic algorithms usually offer ease of development and are much more powerful when customized using knowledge specific to the application. This could be achieved by the use of customized representation (Rogers and Hopfinger, 1994; Rogers, 1996), specialized genetic operators (Venkatasubramanian et al., 1996; Venkata-
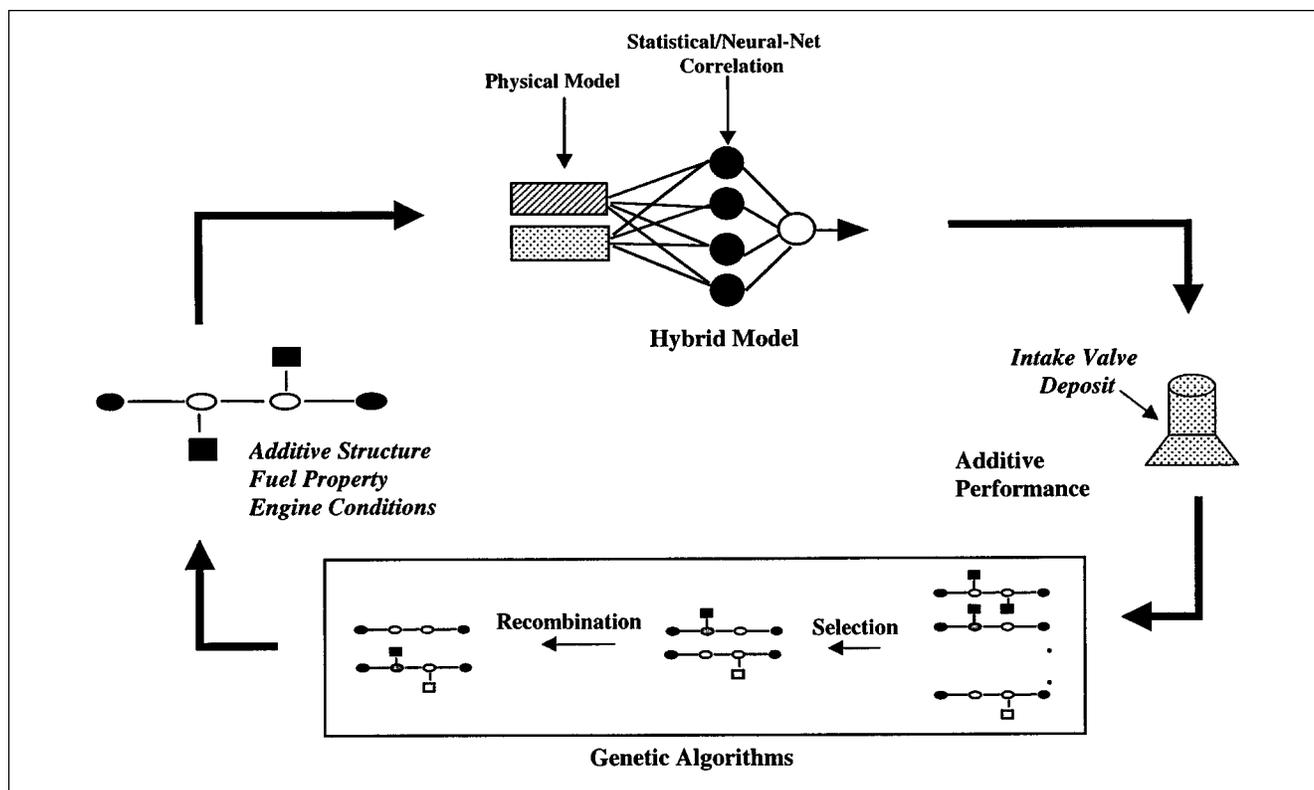
Figure 16. Forward and inverse problems in computer-aided fuel-additive design.
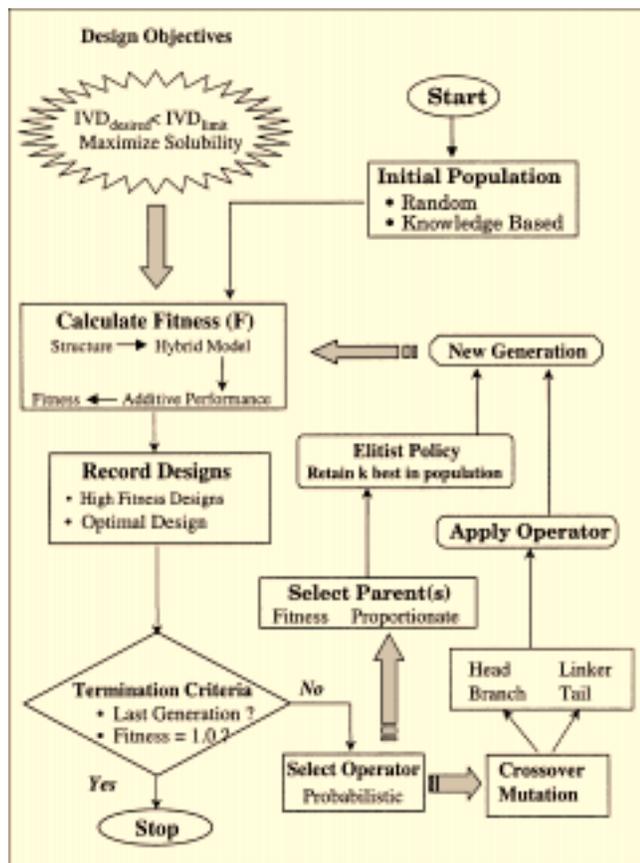
Figure 17. Evolutionary framework for fuel additive design.

Figure 18. Representation used for evolutionary design of fuel additives.

subramanian and Sundaram, 1998; Glen and Payne, 1995), and/or selection schemes (Venkatasubramanian et al., 1995). In this section, the development of an evolutionary algorithm for fuel-additives design is described. The important aspects of this algorithm are the representation and the operators used to reflect the special nature of the product-design problem. The general approach to the fuel-additives design problem is outlined in Figure 16. While the specific details of the product-design problems vary, the main components of the evolutionary algorithm for fuel-additive design were the same as are depicted in Figure 17. The main differences are in the representation, the genetic operators, and in the determination of the fitness function.

Representation

The vital components of the additive that formed a conduit between its structure and performance were the head, linker, and tail. This was in fact also the level of detail at which design decisions were made during formulation. The alphabet set or the basic palette from which additive structures were assembled, therefore, consisted of heads, tails, and linkers. Additional information in terms of satisfying the basic chemical and formulation constraints in the combination of head, linkers, and tails was also incorporated.

Figure 18 shows the representation used in the evolutionary scheme. As opposed to the traditional binary representation (Holland, 1975; Goldberg, 1989; Back and Schwefel,

1993), and the modified string representation (Chan, 1994), the one used here is best described as object oriented. For instance, all head objects contain information on their valency that dictates the number of linkers that should be attached to it. In addition, they also contain details regarding where the branch points are, in terms of the functional groups and kinds of linkages allowed or disallowed. The linkers and the tails are similarly represented. The choice of a particular type of head automatically constrains the number and sometimes the type of the linkers attached to it. A similar relationship would exist between each linker and the tails attached to the linker. This implies that changes to one component of the structure will affect other components connected to it. In this fashion, the object-oriented representation gives greater flexibility in capturing the basic structural knowledge as well as the connectivity underlying the fuel-additive design problem. If an alternative representation such as binary or string was used, it would have necessitated the use of externally imposed constraints or constructs to manage the flow of structural information.

Genetic operators

In evolutionary algorithms, genetic operators are the means of manipulating existing solutions to create new ones that diversify or intensify (Michalewicz, 1996) the structural features being explored for potential optimality. In this sense, genetic operators are powerful in that they provide for both exploration and exploitation of the search space. However, traditional genetic operators are stochastic in nature (DeJong, 1975; Goldberg, 1989) and often need customization to capture the constraints inherent to the application (Michalewicz, 1991, 1995). In an earlier work on polymer design (Chan, 1994), the problem structure was implicitly included in the representation, and no explicit constraining of the mechanics of the genetic operators was necessary. In the fuel-additive problem, however, there are constraints explicitly specified in the formulation that need to be acknowledged by the genetic operators to efficiently search the combinatorial space. These constraints disallow certain combinations of heads and linkers or linkers and tails. In addition, the valency of the head and the linker control the allowed branching for each of them, and hence disallow certain heads or linkers to be considered, for instance, during mutation. These combinatorial constraints were explicitly enforced during genetic operation. Four different crossover and four different mutation operators were used for the fuel-additive design. The essential difference between crossover and mutation was the number of
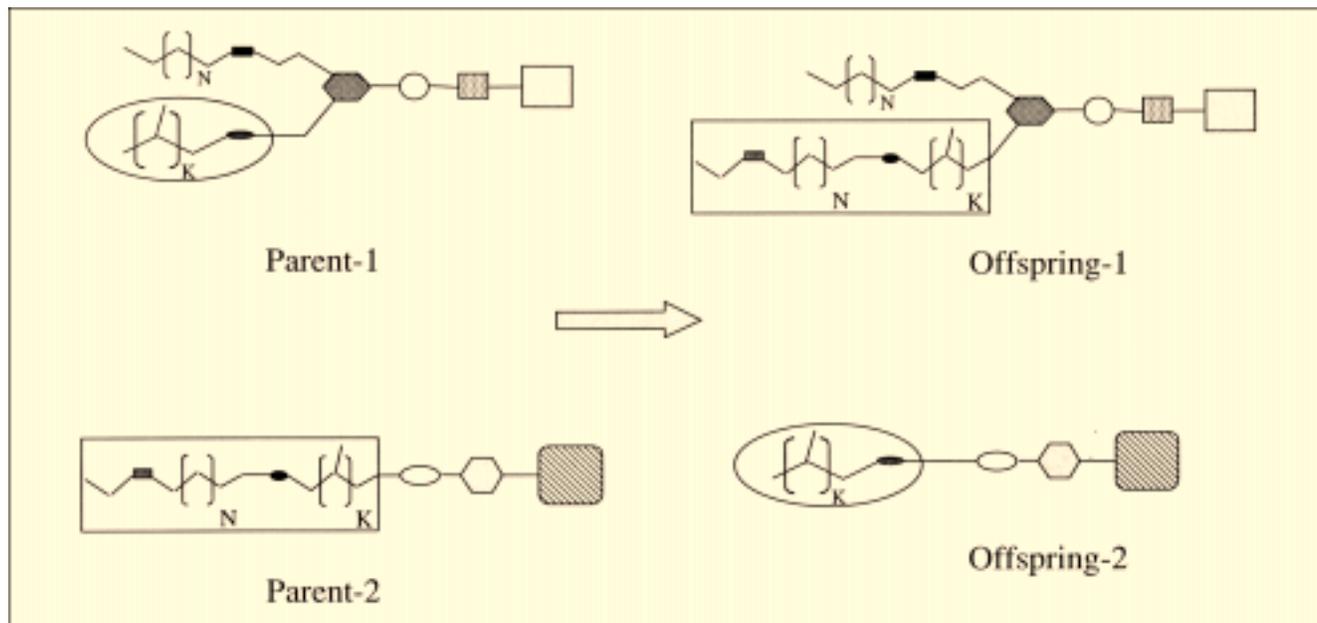
Figure 19. Tail crossover operator.

parents involved in the operation. In applying the genetic operators, the following sequence was followed. An operator was chosen and several attempts made to satisfy feasibility requirements of the anticipated offspring. If all attempts failed (the number of possibilities was always a small finite set), the operation was discarded and a different operator picked for application.

Crossover. Crossover involves exchanges between two parents. The cut points were chosen randomly on each parent and structural features exchanged across the cut points. In
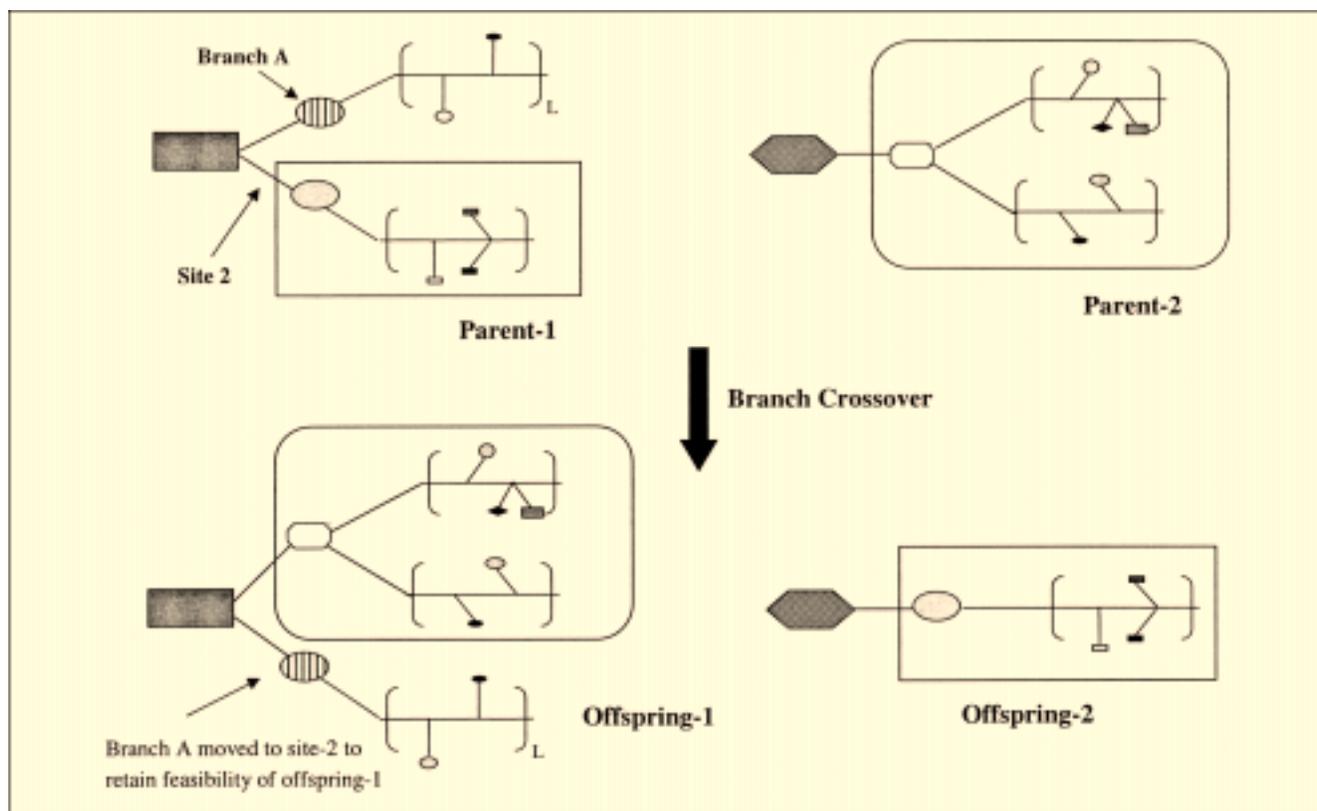
Figure 20. Branch crossover operator.

the fuel-additive structure, however, the cut points could not be chosen at random, as head−linker and linker−tail and tail−tail points provide fixed cut points on the additive structure. Hence, four types of crossover could exist, depending upon the cut point chosen.

Tail Crossover. The linker−tail bond was the cut point for tail crossover. Here, the tails of two parents were exchanged to create two offsprings. This is depicted in Figure 19.

Branch Crossover. The cut point for this operator was the head−linker bond. Here, a linker along with all tails attached to it (a *branch*) was chosen at random on each parent from among all linkers in the respective structures. The chosen branches were then exchanged between the parents to create two offsprings, as shown in Figure 20. Feasibility was enforced, by requiring that the linkers exchanged be compatible with the specific branching sites on the head components of the appropriate parents. For instance, in Figure 20 the crossover was performed only when the branch from the second parent is compatible with attachment site 2 on the head of first parent, and vice versa. In the case of structures with multiple linkers or branches, one of the branches from the unaffected sites on the head may be moved to the open site to ensure that the entering branch can be feasibly placed. As shown in the figure, the first branch of first parent was moved to site 2 to accommodate the entering branch from parent 2 into the first attachment site.

Linker crossover. The linker crossover involved the exchange of chosen compatible linkers between two parents. It is analogous to two-point crossover, which is often used in evolutionary algorithms (Goldberg, 1989; Chan, 1994). The operator was characterized by more than one cut point on each parent, in this case a head−linker attachment site and all the linker−tail attachments of the chosen linker. Such an operator is shown in Figure 21. During application, once two compatible head−linker sites were found on the parents, different rearrangements of tails of the first parent on the second parent's linker and vice versa were explored until all tails on both linkers were compatible to the appropriate linker−tail sites. The linker crossover operation might therefore lead to relocation of some tails to different−tail attachment sites on the new linker.

Head crossover. This operation was similar in flavor to the linker crossover. Head components were exchanged between parents during this operation. The original assignment of linkers to the head might be rearranged as a result of this operator.

Mutation. Mutation involves introduction of localized perturbation to the population through randomly changing one of the structural characteristics of the fuel additive. Four different mutation operators were employed. These were modeled along the same lines as the crossover operators to yield tail, branch, linker, and head crossovers.

Tail Mutation. In the tail-mutation operation, a tail on a chosen parent was replaced with another tail from the base set (that is, set of all possible tails) that was compatible with the linker−tail attachment site from which the original tail was removed.

Branch Mutation. In this operation, an entire branch containing a linker and all tails attached to it were mutated. Once again the linker was chosen randomly, but always from among the set of linkers compatible with the chosen head−linker attachment site. This operator is depicted in Figure 22.
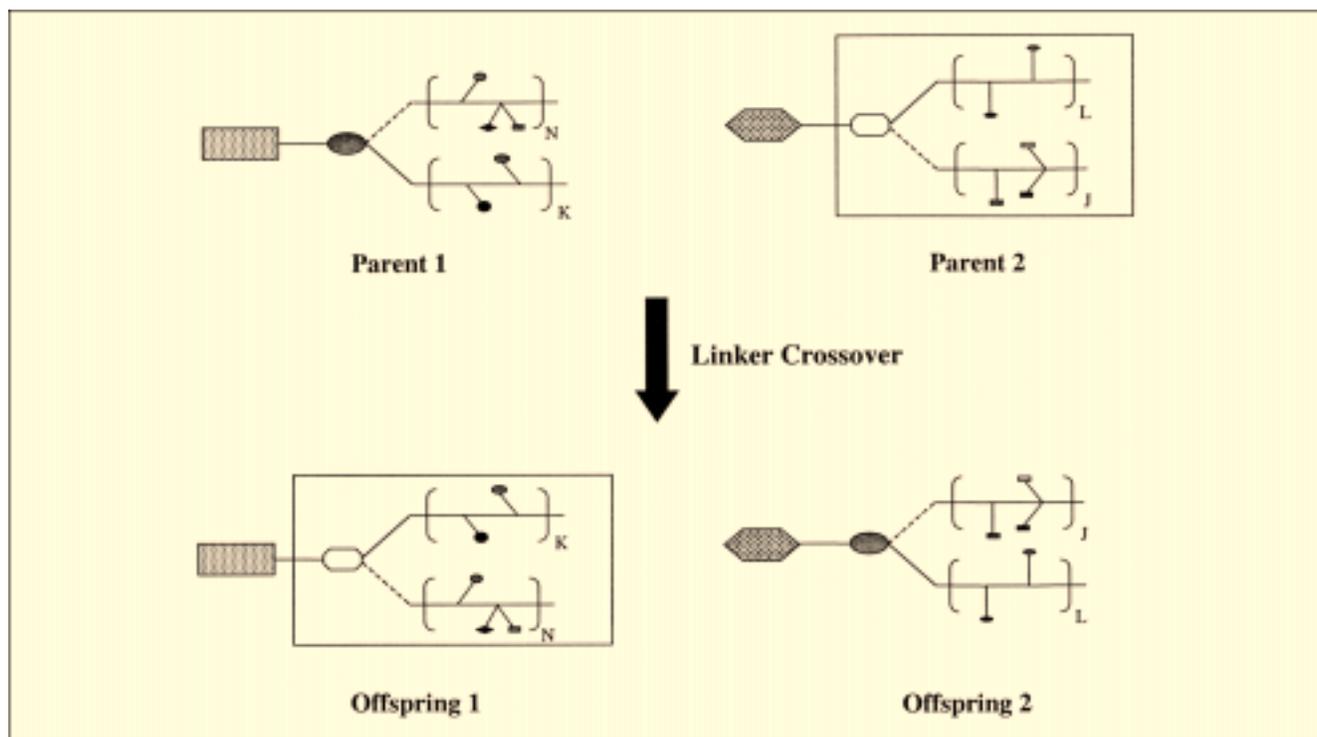
Figure 21. Linker crossover operator.
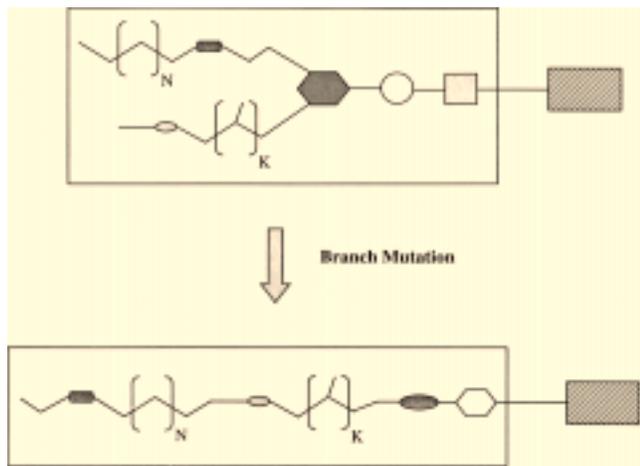
Figure 22. Branch mutation operator.

Linker Mutation. The linker-mutation operation replaced a linker on a chosen parent with a compatible linker from the base set. This was done in two steps. First, a linker was chosen at random from a set of linkers compatible with the head−linker attachment site. This linker was used as the replacement if a sufficient number (as many as the branching of the new linker) of tails on the removed linker were compatible with the sites on the new one. If not, different rearrangements of the tails were explored to identify a compatible configuration. If no such arrangement could be found, this linker was dropped from the list of compatibles (only for this selection procedure) and a different linker component was chosen. If the chosen linker had more branch points than the linker mutated, the additional branches/tails were chosen at random, but compatible with the vacant attachment sites on the linker. If the chosen linker had less branch points, the surplus tails after feasible assignment were discarded.

Head Mutation. In the head-mutation operation, the head of the chosen parent was replaced with a compatible head from the base set. As with the linker crossover, the linkers on the tail might be rearranged as a result of this operation to ensure feasibility of the offspring. All the operators described earlier ensured the creation of feasible offspring from feasible parents. The initial population for the evolutionary algorithm was also created in a random but feasible manner. While none of the genetic operators described earlier were involved at this stage, a feasible growth operator was implemented to assemble the fuel additive, one component at time.

Results and Discussion

As shown in Figure 17, the main aspects of the evolutionary algorithm are the creation of an initial population, fitness determination, selection, and application of genetic operators. In our study, the populations were selected based on fitness proportionate schemes (Goldberg, 1989). The objective function for the search was allowed to be one of two different forms:

1. Maximizing the stability of the additive at a given time and for a given fuel Hildebrand parameter. In this case, the amount of active additive at the desired time and for the given

Table 1. Fitness of Additives Identified for Solubility Maximization

δ (MPa$^{1/2}$)	τ		
	1	5	10
19	0.956	0.791	0.637
21	0.956	0.849	0.729
23	0.923	0.720	0.607
25	0.418	0.406	0.501

fuel Hildebrand parameter was maximized. This objective function did not require the use of the neural-network/regression model.

2. Minimizing the predicted IVD for the given additive structure and operating conditions. For this objective function, a reliably accurate hybrid model (first-principles + NN hybrid) was used for prediction.

The evolutionary algorithm was applied for the design of fuel additives using both the preceding objectives separately. The base set consisted of 25 heads, 9 linkers, and 9 tails. A population of 25 molecules was used in every generation, and the algorithm evolved through 25 generations. The mutation frequency was set at 0.40 and the crossover rate at 0.60. Within mutation, branch-mutation frequency was set at 0.10 and the rest equally shared between the other types. Within crossover, all the different types had the same frequency (0.25). Even for this relatively small-sized problem, the combinatorial size of the search space is around one million. In this study, only the more common heads, tails, and linkers were considered, and hence the base set was rather small. The main aim was to demonstrate the success of the evolutionary approach using constrained operators to come up with high fitness solutions while examining a very small fraction of the total search space. For the IVD prediction objective, the best PLS-NN model based on the Honda database was used. For the solubility objective, the fuel Hildebrand parameter could be varied to reflect the nature of the fuel for which the additive was being designed. Table 1 shows the results from different runs of the evolutionary algorithm using the solubility objective. Each row of the table corresponds to a different fuel Hildebrand parameter setting, the severity of the fuel increasing as indicated across the first column. The second column shows the maximum amount of active additive at time $\tau = 1$ corresponding to the best additive structure located by the evolutionary search. Similarly, columns 3 and 4 show the maximum amount of active additive corresponding to the best structure for the given fuel Hildebrand parameter, but at different times ($\tau = 5, 10$). Since the dimensionless time could be directly correlated to the molecular weight of the additive, the different times are analogous to the activity that could be delivered by the additive structure at a given stage in its degradation. The structures identified to be the best for each set of conditions (τ, δ) were different. This indicated that the choice of the optimal additive structure was dependent on the nature of the fuel itself. This was expected, as the nature of the fuel, and especially its polarity, is crucial in determining the activity of the additive, and hence different additives give different levels of activity in the same fuel. Table 2 shows the solubility performance (in terms of fitness) of the best additive at time $\tau = 1$ as subsequent time is elapsed. As the table shows, the fitness of the additive for

Table 2. Solubility at Longer Times of Additives that are Highly Active at $\tau = 1$

δ (MPa$^{1/2}$)	Solubility (in terms of fitness)	
	$\tau = 5$	$\tau = 10$
19	0.771	0.552
21	0.818	0.690
23	0.671	0.454
25	0.312	0.082

longer times was always lower than that of the best structure identified by the evolutionary scheme at those times.

This demonstrated that the algorithm was able to utilize the degradative stability built into the first-principles model in decision making involving the choice of the structural components. For all these cases, the amount of active additive was estimated by considering a narrow band of Hildebrand values around that of the fuel. This was different from the activity estimates used for IVD prediction. Minimization of the predicted IVD was used as an alternative objective in determining the optimum additive structure. The PLS-NN model developed for the Honda database was used as the predictor. This model did not require the specification of the fuel Hildebrand value, as mentioned earlier. In addition, no time points need to be specified, since these are internal to the IVD prediction model. Once the predicted IVD was determined using the PLS-NN model, the following fitness function was used.

$$F = 1.0; \text{IVD}_{pred} \leq \text{IVD}_{limit}$$

$$F = e^{-(\alpha[\text{IVD}_{prep} - \text{IVD}_{limit}])}; \text{IVD}_{pred} > \text{IVD}_{limit}. \quad (9)$$

The fitness varies between 0 and 1, with the maximum value attained for any value of the predicted IVD less than a preset limit. The fitness decayed exponentially to zero with predicted IVD larger than the limit. The base set and internal parameters for the evolutionary algorithm were the same as the ones used earlier. The limit on the IVD was set at 10 mg, with an assumed initial dosage of 50 PTB. The results of three different runs starting with different initial populations are shown in Table 3. As indicated previously, the proprietary nature of the work does not permit us to reveal the actual

structures in Table 3. The initial population was created as copies of a single randomly assembled dispersant. The figure shows the fitnesses of some of the best structures identified in each run with a description of the structure in terms of major differences from the ones present in the databases. In the table, a novel structure was characterized as a structure where all the components (the head, tails, and linkers) had never occurred in any combination in either the BMW or Honda engine-test databases. Many structures, some of whose components had been examined in combination in the databases, were characterized variants of known structures. Even with a very small sampling of the search space (about 625 out of a possible one million), the evolutionary algorithm was successful in identifying diverse structures that met or were close in meeting the set objectives. Indeed, some of the best structures found in each of the runs were never encountered before. These structures especially III-1 in Table 3, not only have the desired performance measure, but also good synthesis potential.

Conclusions

A hybrid ANN forward model was developed to predict the intake-valve deposit characteristics of fuel additives under given operation conditions. This model was trained on tested on two different sets of engine tests and was proved to have very good predictive accuracy, given the sparseness and noise of the data. The extension of this model from additive structures to blends also showed reasonable promise (using the Ford 2.3L engine-test data) in terms of its ability to capture the best functional aspects of an expert's knowledge. An evolutionary algorithm that designed structures to meet a desired IVD performance was developed. The algorithm was augmented with formulation knowledge with carefully chosen operators that retain structural and functional feasibility of the offspring at every stage. The success of the algorithm was demonstrated with two different objectives and different scenarios. The algorithm was able to identify a whole spectrum of structures, including known good performers and some novel ones. The synthesis potential of the identified structures was ensured to be at least fair, with the use of the feasibility-aware genetic operators.

Table 3. Results from Three Different Runs of the Evolutionary Algorithm

Run	Rank	Fitness	Predicted IVD (hybrid-NN model) (mg)	Structural Description
I	1, I-1	0.997	11.4	Novel Structure. Rarely used components
	2, I-1	0.996	11.5	Novel Structure. Similar to best structure, different core
	6, I-6	0.993	12.0	Variant of structure found in the BMW database
II	1, II-1	0.999	10.1	Novel Structure. Different from 1-I. Infrequently used transitional component
	2, II-2	0.989	12.6	Slight variation of structure found in Honda and BMW databases
	4, II-4	0.983	13.2	Minor variations of structure II-2 earlier
III	1, III-1	1.00	8.9	Novel Structure. Diffusion from 1-I and II-1. Commonly used components
	2, III-2	0.994	11.9	Variation of III-1. One of the transition and branch components different
	3, III-3	0.993	12.1	Variant of structure II-2 earlier. Slight modification of the core. Contains an additional branch

There is another subtle, albeit important use of the hybrid approach developed here. The first-principles model captures the current understanding of the different factors that influence the deposit removal mechanism. The ANN or regression model is the bridge that connects the reality (or the actual IVD performance) to the extent of the current knowledge (the functional descriptors from the first-principles model). As demonstrated earlier, the evolutionary algorithm could be used to optimize either on the descriptors or on the final performance. In addition, the optimization provides a variety of optimal or near-optimal structures for each of these objectives. By studying the differences between the structures based on the best current understanding and those based on the actual performance, some missing aspects in the modeling might become evident. For instance, different kinds of structures that are expected to be optimal in the functional sense may not perform well in terms of the predicted IVD. This information could act as important feedback in different ways. First, they might direct the amplification of engine-test data by inclusion of specific kinds of additives for which the differences were most pronounced. Second, they might serve in refining the first-principles model to account for uncertain parameters, such as measurement of fuel character and so on. Third, they might point toward phenomena that would have to be modeled to account for large differences between theoretical expectations and engine-test data. All these avenues are currently being explored.

The forward model development as well as the inverse problem of fuel-additive design is the first attempt of its kind in this area, to the best of our knowledge. Although the feasibility of reasonably accurate predictions and subsequent design based on these predictor models was proved, more could be done to improve these models. While solubility/stability proved to be a sufficiently good phenomenological descriptor, future work would need to quantify other descriptors, including deposit removal ability and hardware effects. From the perspective of design, future extensions could include total package design, including design of blends and secondary components that enhance fuel-additive performance.

Acknowledgments

Literature Cited

American Society for Testing and Materials (ASTM), *Standard Test Method for Vehicle Evaluation of Unleaded Automotive Spark-Ignition Fuel for Intake Valve Deposit Formation*, Test D5500-94, Philadelphia (1995).

Andersson, G., P. Kaufmann, and L. Renberg, "Nonlinear Modeling with a Coupled Neural Network−PLS Regression System," *J. Chemometrics*, **10**, 605 (1996).

Arters, D. C., E. A. Schiferl, and D. T. Daly, *Variability of Intake Valve Deposit Measurements in the BMW Vehicle Intake Valve Deposit Test*, SAE Tech. Paper Ser., **SP**-1277 (971723), 67 (1997).

Back, T., and H. P. Schwefel, *Evol. Comput.*, **1**, 1 (1993).

Bailey, B. K., L. L. Stavinoha, and D. L. Present, *Gasoline Deposit Mechanism on Aluminium at Various Temperature*, SAE Tech. Paper Ser.: Lubricants and Fuels, **103**:4 (892117): 1432 (1995).

Barton, A. M., *CRC Handbook of Solubility Parameters and Other Cohesion Parameters*, CRC Press, Boca Raton, FL (1991).

Chan, K., *Computer-Aided Molecular Design Using Genetic Algorithms*, PhD Thesis, Purdue University, West Lafayette, IN (1994).

Charalambous, C., "Conjugate Gradient Algorithm for Efficient Training of Artificial Neural-Networks," IEEE Proceedings, **139**(3), 302 (1992).

ChengShi Wai, S., *The Physical Parameters that Influence Deposit Formation on an Intake Valve*, SAE Tech. Paper Ser.: Lubricants and Fuels, **101**:4 (922257), 1404 (1992).

Daneshgari, P., K. Borgmann, and H. Job, *The Influence of Temperature Upon Gasoline Deposit Build-Up on the Intake Valves*, SAE Tech Paper Ser.: Lubricants and Fuels, **98**:4 (890215), 171 (1989).

DeJong, K. A., "An Analysis of a Class of Genetic Adaptive Systems," PhD Thesis, Univ. of Michigan (1975).

Geladi, P., and B. R. Kowalski, "Partial Least-Squares Regression: A Tutorial," *Anal. Chim. Acta*, **185**, 1 (1986).

Gibbs, L. M., *Gasoline Additives— When and Why?* SAE Tech. Paper Ser.: Lubricants and Fuels, **99**:4 (902014), 618 (1990).

Glen, R. C., and A. W. R. Payne, "A Genetic Algorithm for the Automated Generation of Molecules with Constraints," *J. Comput. Aid. Mol. Des.*, **9**, 181 (1995).

Goldberg, D. E., *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, Reading, MA (1989).

Graham, J. P., and B. Evans, *Effect of Intake Valve Deposits on Driveability*, SAE Tech. Paper Ser.: Lubricants and Fuels, **101**:4 (922259), p. 1231 (1992).

Grant, L. J., and R. L. Mason, *SwRI-BMW N.A. Intake Valve Deposit Test— A Statistical Review*, SAE Tech. Paper Ser.: Lubricants and Fuels, **101**:4 (922215), 1221 (1992).

Hagan, M. T., and M. Menhaj, "Training Feedforward Networks with the Marquadt Algorithm," *IEEE Trans. Neural Networks*, **5**(6), 989 (1994).

Haykin, S., *Neural Networks: A Comprehensive Foundation*, Prentice-Hall (1999).

Hildebrand, J. H., and R. L. Scott, *Regular Solutions* (1962).

Holland, J. H., *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor (1975).

Homan, H. S., *The Effect of a Gasoline Additive, Automobile Make and Driving Cycle on Intake Valve Deposits and Combustion Chamber Deposits in a Ten Car Fleet Test*, SAE Tech. Paper Ser.: Lubricants and Fuels, **106**:4 (972836), 863 (1997).

Houser, K. R., and T. A. Crosby, *The Impact of Intake Valve Deposits on Exhaust Emissions*, SAE Tech. Paper Ser.: Lubricants and Fuels, **103**:4 (922259), 1432 (1995).

Kalghatgi, G. T., *Deposit in Gasoline Engines— A Literature Review*, SAE Tech. Paper Ser.: Lubricants and Fuels, **99**:4 (902015), 639 (1990).

Lacey, P. I., K. H. Kohl, L. L. Stavinoha, and R. M. Estefan, *A Laboratory-Scale Test to Predict Intake Valve Deposits*, SAE Tech. Paper Ser.: Lubricants and Fuels, **106**:4 (972833), 880 (1997).

Michalewicz, Z., *Genetic Algorithms+ Data Structure= Evolution Programs* (1996).

Michalewicz, Z., "Genetic Algorithms, Numerical Optimization and Constraints," *Proc. Int. Conf. on Genetic Algorithms*, L. J. Eshelman, ed., San Mateo, CA, p. 151 (1995).

Michalewicz, Z., "Handling Constraints in Genetic Algorithms," *Proc. Int. Conf. on Genetic Algorithms*, R.a.B. Belew, ed., San Mateo, CA, Third ed., Springer-Verlag, Berlin, p. 151 (1991).

Mitsui, J., *Effect of Gasoline Engine Oil Components on Intake Valve Deposit*, SAE Tech. Paper Ser.: Lubricants and Fuels, **102**:4 (932792), 1714 (1993).

Qin, S. J., and T. J. McAvoy, "Nonlinear PLS Modeling Using Neural Networks," *Comput. Chem. Eng.*, **16**, 379 (1992).

Rogers, D., and A. J. Hopfinger, "Application of Genetic Function Approximation to Quantitative Structure-Activity Relationships and Quantitative Structure-Property Relationships," *J. Chem. Inf. Comput. Sci.*, **34**, 854 (1994).

Rogers, D., "Some Theory and Examples of Genetic Function Approximation with Comparison to Evolutionary Techniques," *Genetic Algorithms in Molecular Modeling*, J. Devillers, ed., London, p. 87 (1996).

Rummelhart, D. E., G. E. Hinton, and R. J. Williams, "Learning Representation of Back-Propagation Errors," *Nature (London)*, **323**, 533 (1986).

Schenker, B., and M. Agarwal, ''Cross-Validated Structure Selection for Neural-Networks,'' *Comput. Chem. Eng.*, **20**, 175 (1996).

U. S. Environmental Protection Agency (EPA), *Final Rule: Certification Standards for Deposit Control Gasoline Additives*, Resources, Office of Management, Washington, DC (1996).

Venkatasubramanian, V., and A. Sundaram, ''Genetic Algorithms: Introduction and Applications,'' *Encyclopedia of Computational Chemistry*, P.V.R. Schleyer, N. L. Allinger, T. Clark, J. Gasteiger, P. A. Kollman, H. F. Schaefer III, and P. R. Schreiner, eds., Wiley, Chichester, U.K., pp. 1115−1135 (1998).

Venkatasubramanian, V., K. Chan, and J. M. Caruthers, ''Genetic Algorithmic Approach for Computer-Aided Molecular Design,'' *Computer-Aided Molecular Design*, p. 396 (1995).

Venkatasubramanian, V., K. Chan, A. Sundaram, and J. M. Caruthers, ''Computer-aided Molecular Design Using Neural-Networks and Genetic Algorithms,'' *Genetic Algorithms in Molecular Modeling*, J. Devillers, ed., London, p. 271 (1996).