

Hybrid Soft Computing Systems: Industrial and Commercial Applications

PIERO P. BONISSONE, MEMBER, IEEE, YU-TO CHEN, MEMBER, IEEE,
KAI GOEBEL, AND PRATAP S. KHEDKAR

Invited Paper

Soft computing (SC) is an association of computing methodologies that includes as its principal members fuzzy logic, neurocomputing, evolutionary computing and probabilistic computing. We present a collection of methods and tools that can be used to perform diagnostics, estimation, and control. These tools are a great match for real-world applications that are characterized by imprecise, uncertain data and incomplete domain knowledge. We outline the advantages of applying SC techniques and in particular the synergy derived from the use of hybrid SC systems. We illustrate some combinations of hybrid SC systems, such as fuzzy logic controllers (FLC's) tuned by neural networks (NN's) and evolutionary computing (EC), NN's tuned by EC or FLC's, and EC controlled by FLC's. We discuss three successful real-world examples of SC applications to industrial equipment diagnostics, freight train control, and residential property valuation.

Keywords— *Computational intelligence, control systems, evolutionary computation, fuzzy systems, neural networks, property valuation, soft computing.*

I. INTRODUCTION

A. Motivation

There is a wide gamut of industrial and commercial problems that requires the analysis of uncertain and imprecise information. Usually, an incomplete understanding of the problem domain further compounds the problem of generating models used to explain past behaviors or predict future ones. These problems present a great opportunity for the application of soft computing (SC) technologies.

In the industrial world, it is increasingly common for companies to provide diagnostics and prognostics services for expensive machinery (e.g., locomotives, aircraft engines, medical equipment). Many manufacturing companies are trying to shift their operations to the service field where they expect to find higher margins. One embodiment of this

trend is the successful offering of long-term service agreements with guaranteed up time. These contracts provide economic incentives for the service provider to keep the equipment in working order. Since performing unnecessary maintenance actions is undesired and costly, in the interest of maximizing the margins, service should only be performed when required. This can be accomplished by using a tool that measures the state of the system and indicates any incipient failures. Such a tool must have a high level of sophistication that incorporates monitoring, fault detection, decision making about possible preventive or corrective action, and monitoring of its execution. A second industrial challenge is to provide intelligent automated controllers for complex dynamic systems that are currently controlled by human operators. For instance, automatic freight train handling is a task that still eludes full automation.

In the commercial and financial world, we can still find inefficient markets in which timely asset valuation and pricing still represent a big challenge. For instance, while we can easily determine the value of a portfolio of stocks and bonds, we cannot perform the same precise valuation for a portfolio of real-estate assets, which could have been used as collateral for mortgage-backed securities.

Due to of the complexity of these tasks, artificial intelligence (AI), and in particular SC, have been called upon to help. Recently we have applied SC techniques in support of service tasks such as anomaly detection and identification, diagnostics, prognostics, estimation and control [1], [2], and residential property valuation [3]. This paper focuses on the use of SC on these three aspects: diagnostics; control; and valuation. To begin, we will present our interpretation of the definition and scope of the concept of SC.

B. SC

SC is a term originally coined by Zadeh [4] to denote systems that "... exploit the tolerance for imprecision, uncertainty, and partial truth to achieve tractability, robustness, low solution cost, and better rapport with reality." SC

Manuscript received March 29, 1999; revised June 3, 1999.

The authors are with the General Electric Corporate Research and Development Center, Niskayuna, NY 12309 USA.

Publisher Item Identifier S 0018-9219(99)06916-9.

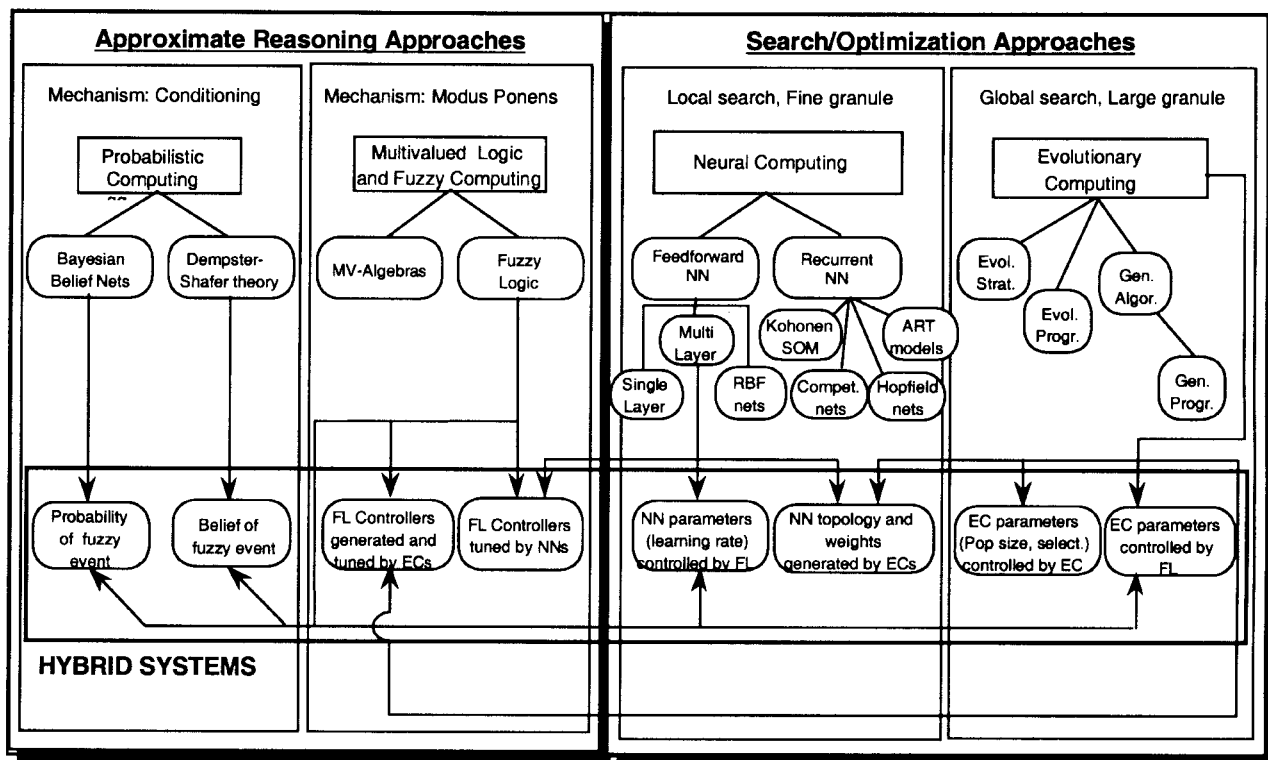


Fig. 1. SC components and hybrid systems.

is “an association of computing methodologies that includes as its principal members fuzzy logic (FL), neuro-computing (NC), evolutionary computing (EC) and probabilistic computing (PC)” [5]. However, it should be noted that we have not reached a consensus yet as to the exact scope or definition of SC [6].

The main reason for the popularity of SC is the synergy derived from its components. SC’s main characteristic is its intrinsic capability to create hybrid systems that are based on a (loose or tight) integration of constituent technologies. This integration provides complementary reasoning and searching methods that allow us to combine domain knowledge and empirical data to develop flexible computing tools and solve complex problems. Fig. 1 illustrates a taxonomy of these hybrid algorithms and their components. Extensive coverage of this topic can be found in [7] and [8].

C. SC Components and Taxonomy

1) *Fuzzy Computing*: The treatment of imprecision and vagueness can be traced back to the work of Post, Kleene, and Lukasiewicz, multiple-valued logicians who in the early 1930’s proposed the use of three-valued logic systems (later followed by infinite-valued logic) to represent undetermined, unknown, or other possible intermediate truth-values between the classical Boolean true and false values [9]. In 1937, the philosopher Black suggested the use of a consistency profile to represent vague concepts [10]. While vagueness relates to ambiguity, fuzziness addresses the lack of sharp set boundaries. It was not until 1965, when Zadeh proposed a complete theory of fuzzy sets (and its

isomorphic fuzzy logic), that we were able to represent and manipulate ill-defined concepts [11].

In a narrow sense, FL could be considered a fuzzification of Lukasiewicz Aleph-1 multiple-valued logic [12]. In the broader sense, however, this narrow interpretation represents only one of FL’s four facets [13]. More specifically, FL has a logical facet, derived from its multiple-valued logic genealogy, a set-theoretic facet, stemming from the representation of sets with ill-defined boundaries, a relational facet, focused on the representation and use of fuzzy relations, and an epistemic facet, covering the use of FL to fuzzy knowledge-based systems and databases. A comprehensive review of FL and fuzzy computing can be found in [14].

FL gives us a language, with syntax and local semantics, in which we can translate qualitative knowledge about the problem to be solved. In particular, FL allows us to use linguistic variables to model dynamic systems. These variables take fuzzy values that are characterized by a label (a sentence generated from the syntax) and a meaning (a membership function determined by a local semantic procedure). The meaning of a linguistic variable may be interpreted as an elastic constraint on its value. These constraints are propagated by fuzzy inference operations, based on the *generalized modus-ponens*. This reasoning mechanism, with its interpolation properties, gives FL a robustness with respect to variations in the system’s parameters, disturbances, etc., which is one of FL’s main characteristics [15].

2) *PC*: Rather than retracing the history of probability, we will focus on the development of PC and illustrate

the way it complements fuzzy computing. As depicted in Fig. 1, we can divide probabilistic computing into two classes: single-valued and interval-valued systems.

Bayesian belief networks (BBN's), based on the original work of Bayes [16], are a typical example of single-valued probabilistic reasoning systems. They started with approximate methods used in first-generation expert systems, such as MYCIN's confirmation theory [17] and PROSPECTOR's modified Bayesian rule [18], and evolved into formal methods for propagating probability values over networks [19]–[20]. In general, probabilistic reasoning systems have exponential complexity, when we need to compute the joint probability distributions for all the variables used in a model. Before the advent of BBN's, it was customary to avoid such computational problems by making unrealistic, global assumptions of conditional independence. By using BBN's we can decrease this complexity by encoding domain knowledge as structural information: the presence or lack of conditional dependency between two variables is indicated by the presence or lack of a link connecting the nodes representing such variables in the network topology. For specialized topologies (trees, polytrees, directed acyclic graphs), efficient propagation algorithms have been proposed by Kim and Pearl [21]. However, the complexity of multiple-connected BBN's is still exponential in the number of nodes of the largest subgraph. When a graph decomposition is not possible, we resort to approximate methods, such as clustering and bounding conditioning, and simulation techniques, such as logic samplings and Markov simulations.

Dempster–Shafer (DS) systems are a typical example of interval-valued probabilistic reasoning systems. They provide lower and upper probability bounds instead of a single value as in most BBN cases. The DS theory was developed independently by Dempster [22] and Shafer [23]. Dempster proposed a calculus for dealing with interval-valued probabilities induced by multiple-valued mappings. Shafer, on the other hand, started from an axiomatic approach and defined a calculus of belief functions. His purpose was to compute the credibility (degree of belief) of statements made by different sources, taking into account the sources' reliability. Although they started from different semantics, both calculi were identical.

PC provides a way to evaluate the outcome of systems affected by randomness (or other types of probabilistic uncertainty). PC's basic inferential mechanism—conditioning—allows us to modify previous estimates of the system's outcome based on new evidence.

3) *Comparison Between PC and Fuzzy Computing:* In this brief review of fuzzy computing and PC, we would like to emphasize that randomness and fuzziness capture two different types of uncertainty. In randomness, the uncertainty is derived from the nondeterministic membership of a point from a sample space (describing the set of possible values for the random variable), into a well-defined region of that space (describing the event). A probability value describes the tendency or frequency with which the random variable takes values inside the region. In fuzziness, the uncertainty

is derived from the deterministic but partial membership of a point (from a reference space) into an imprecisely defined region of that space. The region is represented by a fuzzy set. The characteristic function of the fuzzy set maps every point from such space into the real-valued interval $[0,1]$, instead of the set $\{0,1\}$. A partial membership value does not represent a frequency. Rather, it describes the degree to which that particular element of the universe of discourse satisfies the property that characterizes the fuzzy set. In 1968, Zadeh noted the complementary nature of these two concepts, when he introduced the probability measure of a fuzzy event [24]. In 1981, Smets extended the theory of belief functions to fuzzy sets by defining the belief of a fuzzy event [25]. These are the first two cases of hybrid systems illustrated in Fig. 1.

4) *Neural Computing:* The genealogy of neural networks (NN's) goes back to 1943, when McCulloch and Pitts showed that a network of binary decision units (BDN's) could implement any logical function [26]. Building upon this concept, Rosenblatt proposed a one-layer feedforward network, called a perceptron, and demonstrated that it could be trained to classify patterns [27]–[29]. Minsky and Papert [30] proved that single-layer perceptrons could only provide linear partitions of the decision space. As such they were not capable of separating nonlinear or nonconvex regions. This caused the NN community to focus its efforts on the development of multilayer NN's that could overcome these limitations. The training of these networks, however, was still problematic. Finally, the introduction of backpropagation (BP), independently developed by Werbos [31], Parker [32], and LeCun [33], provided a sound theoretical way to train multilayered, feedforward networks with nonlinear activation functions. In 1989, Hornik *et al.* proved that a three-layer NN (with one input layer, one hidden layer of squashing units, and one output layer of linear units) was a universal functional approximator [34].

Topologically, NN's are divided into feedforward and recurrent networks. The feedforward networks include single- and multiple-layer perceptrons, as well as radial basis functions (RBF) networks [35]. The recurrent networks cover competitive networks, self-organizing maps (SOM's) [36], Hopfield nets [37], and adaptive resonance theory (ART) models [38]. While feedforward NN's are used in supervised mode, recurrent NN's are typically geared toward unsupervised learning, associative memory, and self organization. In the context of this paper, we will only consider feedforward NN's. Given the functional equivalence already proven between RBF and fuzzy systems [39] we will further limit our discussion to multilayer feedforward networks. A comprehensive current review of neurocomputing can be found in [40].

Feedforward multilayer NN's are computational structures that can be trained to learn patterns from examples. They are composed of a network of processing units or neurons. Each neuron performs a weighted sum of its input, using the resulting sum as the argument of a nonlinear activation function. Originally the activation

functions were sharp thresholds (or Heavyside) functions, which evolved to piecewise linear saturation functions, to differentiable saturation functions (or sigmoids), and to Gaussian functions (for RBF's). By using a training set that samples the relation between inputs and outputs, and a learning method that trains their weight vector to minimize a quadratic error function, NN's offer the capabilities of a supervised learning algorithm that performs fine-granule local optimization.

5) *EC*: EC algorithms exhibit an adaptive behavior that allows them to handle nonlinear, high dimensional problems without requiring differentiability or explicit knowledge of the problem structure. As a result, these algorithms are very robust to time-varying behavior, even though they may exhibit low speed of convergence. EC covers many important families of stochastic algorithms, including evolutionary strategies (ES), proposed by Rechenberg [41] and Schwefel [42], evolutionary programming (EP), introduced by Fogel [43]–[44], and genetic algorithms (GA's), based on the work of Fraser [45], Bremermann [46], Reed *et al.* [47], and Holland [48]–[50], which contain as a subset genetic programming (GP), introduced by Koza [51].

The history of EC is too complex to be completely summarized in a few paragraphs. It could be traced back to Friedberg [52], who studied the evolution of a learning machine capable of computing a given input–output function; Fraser [45] and Bremermann [46], who investigated some concepts of genetic algorithms using a binary encoding of the genotype; Barricelli [53], who performed some numerical simulation of evolutionary processes; and Reed *et al.* [47], who explored similar concepts in a simplified poker game simulation. The interested reader is referred to [54] for a comprehensive overview of evolutionary computing and to [55] for an encyclopedic treatment of the same subject. A collection of selected papers illustrating the history of EC can be found in [56].

We will provide a brief description of four components of EC: ES; EP; GA's; and GP. Evolutionary strategies were originally proposed for the optimization of continuous parameter spaces. In typical ES notation the symbols (μ, λ) -ES and $(\mu + \lambda)$ -ES indicate the algorithms in which a population of μ parents generate λ offspring; the best μ individuals are selected and kept in the next generation. In the case of (μ, λ) -ES the parents are excluded from the selection and cannot be part of the next generation, while in the $(\mu + \lambda)$ -ES they can. In their early versions, ES used a continuous representation and mutation operator working on a single individual to create a new individual, i.e., $(1 + 1)$ -ES. Later, a recombination operator was added to the mutation operator as ES were extended to evolve a population of individuals, as in $(\mu + 1)$ -ES [57] and in $(\mu + \lambda)$ -ES [58]. Each individual has the same opportunity to generate an offspring. Furthermore, each individual has its own mutation operator, which is inherited and may be different for each parameter.

In its early versions, EP was applied to identification, sequence prediction, automatic control, pattern recognition, and optimal gaming strategies [43]. To achieve these goals,

EP evolved finite state machines that tried to maximize a payoff function. However, recent versions of EP have focused on continuous parameter optimization and the training of NN's [59]. As a result, EP has become more similar to ES. The main distinction between EP and ES is that EP focuses on the behavior of species, while ES focus on the behavior of individuals. Therefore, EP does not usually employ crossover as a form of variation operator. EP can be considered as a special case of $(\mu + \mu)$ -ES, in which mutation is the only operator used to search for new solutions.

GA's perform a randomized global search in a solution space by using a genotypic rather than phenotypic representation. Canonical GA's encode each candidate solution to a given problem as a binary, integer, or real-valued string, referred to as chromosome. Each string's element represents a particular feature in the solution. The string is evaluated by a fitness function to determine the solution's quality: better-fit solutions survive and produce offspring, while less-fit solutions are culled from the population. To evolve current solutions, each string can be modified by two operators: crossover and mutation. Crossover acts to capture the best features of two parents and pass it to an offspring. Mutation is a probabilistic operator that tries to introduce needed features in populations of solutions that lack such features. The new individuals created by these operators form the next generation of potential solutions. The problem constraints are typically modeled by penalties in the fitness function or encoded directly in the solution data structures [60].

GP is a particular form of GA in which the chromosome has a hierarchical rather than a linear structure whose size is not predefined. The individuals in the population are tree structured programs, and the genetic operators are applied to their branches (subtrees) or to single nodes. This type of chromosome can represent the parse tree of an executable expression, such as an S-expression in LISP, an arithmetic expression, etc. GP has been used in time-series predictions, control, optimization of NN topologies, etc.

As noted by Fogel [54], ES, EP, and GA's share many common traits: "...Each maintains a population of trial solutions, imposes random changes to those solutions, and incorporates selection to determine which solutions to maintain in future generations..." Fogel also notes that "...GA's emphasize models of genetic operators as observed in nature, such as crossing-over, inversion, and point mutation, and apply these to abstracted chromosomes..." while ES and EP "...emphasize mutational transformations that maintain behavioral linkage between each parent and its offspring."

Finally, we would like to remark that EC components have increasingly shared their typical traits; ES have added recombination operators similar to GA's, while GA's have been extended by the use of real-number-encoded chromosomes, adaptive mutation rates, and additive mutation operators. For the sake of brevity, in the context of this paper, we will limit most of our discussion to GA's, but we encourage the reader to consider evolutionary algorithms broadly.

6) *SC Taxonomy*: The common denominator of these technologies is their departure from classical reasoning and modeling approaches that are usually based on Boolean logic, analytical models, crisp classifications, and deterministic search. In ideal problem formulations, the systems to be modeled or controlled are described by complete and precise information. In these cases, formal reasoning systems, such as theorem provers, can be used to attach binary truth-values to statements that describe the state or behavior of the physical system.

When we solve real-world problems, we realize that such systems are typically ill-defined, difficult to model, and possess large solution spaces. In these cases, precise models are impractical, too expensive, or nonexistent. Our solution must be generated by leveraging two kinds of resources: problem domain knowledge of the process or product and field data that characterize the behavior of the system. The relevant available domain knowledge is typically a combination of first principles and empirical knowledge and is usually incomplete and sometimes erroneous. The available data are typically a collection of input–output measurements, representing instances of the system’s behavior, and may be incomplete and noisy.

We can observe from Fig. 1 that the two main approaches in SC are knowledge-driven reasoning systems (such as probabilistic and fuzzy computing) and data-driven search and optimization approaches (such as neurocomputing and EC). This taxonomy, however, is soft in nature, given the existence of many hybrid systems that span across more than one field.

D. Alternative Approaches to SC

The alternative approaches to SC are the traditional knowledge-driven reasoning systems and the data-driven systems. The first class of approaches are exemplified by first-principle-derived models (based on differential or difference equations), by first-principle-qualitative models (based on symbolic, qualitative calculi [61]–[62]), by classical Boolean systems, such as theorem provers (based on unification and resolution mechanisms), or by expert systems embodying empirical or experiential knowledge. All these approaches are characterized by the encoding of problem domain knowledge into a model that tries to replicate the system’s behavior. The second class of approaches are the regression models and crisp clustering techniques that attempt to derive models from any information available from (or usually buried in) the data.

Knowledge-driven systems, however, have limitations, as their underlying knowledge is usually incomplete. Sometimes, these systems require the use of simplifying assumptions to keep the problem tractable (e.g., linearization, hierarchy of local models, use of default values). Theoretically derived knowledge may even be inconsistent with the real system’s behavior. Experiential knowledge, on the other hand, could be static, represented by a collection of instances of relationships among the system variables (sometimes pointing to causality, more often just highlighting correlation). The result is the creation of precise but

simplified models that do not properly reflect reality, or the creation of approximate models that tend to become stale with time and are difficult to maintain.

Purely data-driven methods also have their drawbacks, since data tend to be high dimensional, noisy, incomplete (e.g., databases with empty fields in their records), or wrong (e.g., outliers due to malfunctioning or failing sensors, transmission problems, erroneous manual data entries). Some techniques have been developed to address these problems, such as feature extraction, filtering and validation gates, imputation models, and virtual sensors that model the recorded data as a function of other variables.

The fundamental problem of these classical approaches lies in representing and integrating uncertain, imprecise knowledge in data-driven methods or in making use of somewhat unreliable data in a knowledge-driven approach.

E. SC Solutions

Although it would be presumptuous to claim that SC solves this problem, it is reasonable to affirm that SC provides a different paradigm in terms of representation and methodologies, which facilitates these integration attempts. For instance, in classical control theory the problem of developing models is decomposed into system identification and parameter estimation. Usually the former is used to determine the order of the differential equations and the latter determines its coefficients. Hence, in this traditional approach we have: model = structure + parameters. This equation does not change with the advent of SC. However, we now have a much richer repertoire to represent the structure, to tune the parameters, and to iterate this process. It is understood that the search method used to find the parameter values is an important and implicit part of the above equation, which needs to be chosen carefully for efficient model construction.

For example, the knowledge base (KB) in a Mamdani-type fuzzy system [63] is typically used to approximate a relationship between a state X and an output Y . The KB is completely defined by: a set of scaling factors (SF’s), determining the ranges of values for the state and output variables; a termset (TS), defining the membership function of the values taken by each state and output variable; and by a ruleset (RS), characterizing a syntactic mapping of symbols from X to Y . The structure of the underlying model is the RS, while the model parameters are the SF’s and TS’s. The inference obtained from such a system is the result of interpolating among the outputs of all relevant rules. The inference’s outcome is a membership function defined on the output space, which is then aggregated (defuzzified) to produce a crisp output. With this inference mechanism we can define a deterministic mapping between each point in the state space and its corresponding output. Therefore, we can now equate a fuzzy KB to a response surface in the cross product of state and output spaces, which approximates the original relationship.

A Takagi–Sugeno–Kang (TSK) type of fuzzy system [64] increases its representational power by allowing the use of a first-order polynomial, defined on the state space, to be the

output of each rule in the ruleset. This enhanced representational power, at the expense of local legibility [65], results in a model that is equivalent to RBF's [66]. The same model can be translated into a structured network, such as the adaptive neural fuzzy inference systems (ANFIS) proposed by Jang [67]. In ANFIS the ruleset determines the topology of the net (model structure), while dedicated nodes in the corresponding layers of the net (model parameters) define the TS's and the polynomial coefficients. Similarly, in the traditional NN's the topology represents the model structure and the links' weights represent the model parameters.

While NN and structured nets use local search methods, such as BP, to tune their parameters, it is possible to use evolutionary-computation-based global search methods to achieve the same parametric tuning or to postulate new structures. An extensive coverage of these approaches can be found in [8].

F. Paper Structure

In Section II we will focus on SC hybrid algorithms and the synergy derived from their use of complementary components. After a brief discussion of the problem of diagnosis, estimation, and control provided in Section III, we will illustrate three applications of SC techniques. The first application, described in Section IV, consists in the adaptation of fuzzy clustering to classify gas turbine anomalies. The second application, illustrated in Section V, covers the use of GA's to tune a fuzzy system that implements an automatic train handler controller. The third application, discussed in Sections VI and VII, describes the use of a customized neural-fuzzy system and a fusion of models to determine the market value of residential properties. In the last section, we summarize the advantages of using SC hybrid systems and discuss some potential extensions of these technologies.

II. HYBRID ALGORITHMS: SYMBIOSIS

Over the past few years we have seen an increasing number of hybrid algorithms, in which two or more SC technologies have been integrated to leverage the advantages of individual approaches. By combining smoothness and embedded empirical qualitative knowledge with adaptability and general learning ability, these hybrid systems improve the overall algorithm performance. We will now analyze a few of such combinations, as depicted in the lower box of Fig. 1. First we will illustrate the application of NN's and EC to tune FL systems, as well as the use of EC to generate and tune NN's. Then, we will see how fuzzy systems can control the learning of NN's and the run-time performance of EC.

A. FL Tuned by NN

ANFIS [67] is a representative hybrid system in which NN's are used to tune a fuzzy logic controller (FLC). ANFIS consists of a six-layer generalized network. The first and sixth layers correspond to the system inputs and outputs. The second layer defines the fuzzy partitions

(TS's) on the input space, while the third layer performs a differentiable T-norm operation, such as the product or the soft minimum. The fourth layer normalizes the evaluation of the left-hand side of each rule, so that their degrees of applicability will add up to one. The fifth layer computes the polynomial coefficients in the right-hand side of each Takagi-Sugeno rule. Jang's approach is based on a two-stroke optimization process. During the forward stroke the TS's of the second layer are kept equal to their previous iteration value while the coefficients of the fifth layer are computed using a least mean square method. At this point ANFIS produces an output that is compared with the one from the training set to produce an error. The error gradient information is then used in the backward stroke to modify the fuzzy partitions of the second layer. This process is continued until convergence is reached.

B. FL Tuned by EC

As part of the design of a fuzzy controller, it is relatively easy to specify a cost function to evaluate the state trajectory of the closed-loop system. On the other hand, for each state vector instance, it is difficult to specify its corresponding desired controller's actions, as would be required by supervised learning methods. Thus evolutionary algorithms can evolve a fuzzy controller by using a fitness function based on such a cost function.

Many researchers have explored the use of EC to tune FLC's. Cordon *et al.* [68] alone list a bibliography of over 300 papers combining GA's with FL, of which at least half are specific to the tuning and design of fuzzy controllers by GA's. For the sake of brevity we will limit this section to a few contributions. These methods differ mostly in the order or the selection of the various FLC components that are tuned (TS's, rules, SF's).

Karr, one of the precursors in this quest, used GA's to modify the membership functions in the TS's of the variables used by the FLC [69]. Karr used a binary encoding to represent three parameters defining a membership value in each TS. The binary chromosome was the concatenation of all TS's. The fitness function was a quadratic error calculated for four randomly chosen initial conditions.

Lee and Takagi tuned both rules and TS's [70]. They used a binary encoding for each three-tuple characterizing a triangular membership distribution. Each chromosome represents a TSK rule, concatenating the membership distributions in the rule antecedent with the polynomial coefficients of the consequent. Most of the above approaches used the sum of quadratic errors as a fitness function. Surmann *et al.* extended this quadratic function by adding to it an entropy term describing the number of activated rules [71].

Kinzel *et al.* departed from the string representation and used a cross-product matrix to encode the rule set (as if it were in table form) [72]. They also proposed customized (point-radius) crossover operators that were similar to the two-point crossover for string encoding. They first initialized the rule base according to intuitive heuristics, used GA's to generate a better rule base, and finally tuned the membership functions of the best rule base.

This order of the tuning process is similar to that typically used by self-organizing controllers [73].

Bonissone *et al.* [74] followed the tuning order suggested by Zheng [75] for manual tuning. They began with macroscopic effects by tuning the FLC state and control variable SF's while using a standard uniformly spread TS and a homogeneous rule base. After obtaining the best scaling factors, they proceeded to tune the TS's, causing medium-size effects. Finally, if additional improvements were needed, they tuned the rule base to achieve microscopic effects. Parameter sensitivity order can be easily understood if we visualize a homogeneous rule base as a rule table. A modified scaling factor affects the entire rule table; a modified term in a TS affects one row, column, or diagonal in the table; a modified rule only affects one table cell.

This approach exemplifies the synergy of SC technologies, as it was used in the development of a fuzzy PI-control to synthesize a freight train handling controller. This application is described in greater detail in Section V.

C. NN's Generated by EC

EC has been successfully applied to the synthesis and tuning of neural networks. Specifically, EC has been used to:

- 1) find the optimal set of weights for a given topology, thus replacing BP;
- 2) evolve the network topology (number of hidden layers, hidden nodes, and number of links), while using BP to tune the weights;
- 3) simultaneously evolve both weights and topology;
- 4) evolve the reward function, making it adaptive;
- 5) select the critical input features for the NN.

An extensive review of these five application areas can be found in [76]–[78]. Here, we will briefly cover three early applications of EC to NN's, focusing on the evolution of the network's weights and topology. BP, the typical local tuning algorithm for NN's, is a gradient-based method that minimizes the sum of squared errors between ideal and computed outputs. BP has an efficient convergence rate on convex error surfaces. However, when the error surface exhibits plateaus, spikes, saddle points, and other multimodal features, BP, like other gradient-based methods, can become trapped in suboptimal regions. Instead of repeating the local search from many different initial conditions, or applying perturbations on the weights when the search seems to stagnate, it is desirable to rely on a more robust, global search method. In 1988, Jurick provided a critique of BP and was among the firsts to suggest the use of robust tuning techniques for NN's, including evolutionary approaches [79].

1) *Using GA's:* Montana and Davis proposed the use of genetic algorithms to train a feedforward NN for a given topology [80]. They used real-valued steady-state GA's with direct encoding and additive mutation. GA's perform efficient coarse-granularity search (finding the promising

region where the global minimum is located) but they may be inefficient in the fine-granularity search (finding the global minimum). These characteristics motivated Kitano to propose a hybrid algorithm in which the GA's would find a good parameter region that was used to initialize the NN [81]. At that point, BP would perform the final parameter tuning. McInerney improved Kitano's algorithm by using the GA to escape from the local minima found by the BP during the training of the NN's (rather than initializing the NN's using the GA's and then tuning it using BP) [82]. They also provided a dynamic adaptation of the NN learning rate.

The simultaneous evolution of the network's weights and topology represents a more challenging task for the GA's. Critical to this task's success is the appropriate selection of the genotypic representation. There are three available methods: direct; parametrized; and grammar encoding.

Direct binary encoding is often implemented by encoding the network as a connectivity matrix of size $N \times N$, where N is the maximum number of neurons in the network. The encoding of the weights requires predetermining the granularity needed to represent the weights. Typically, direct binary encoding results in long chromosomes that tend to decrease the GA's efficiency. To address this problem, Maniezzo [83] suggested the use of variable granularity to represent the weights. Specifically, he proposed using different-length genotypes that included a control parameter, the coding granularity. He also recognized the disruptive effect of the crossover operation and decreased the probability of crossover, relying more on the mutation operator for the generation of new individuals.

Parametrized encoding uses a compact representation for the network, indicating the number of layers and the number of neurons in each layer, instead of the connectivity matrix [84]–[85]. However, this type of encoding can only represent layered architectures, e.g., fully connected networks without any connection from input to output layers.

Grammar encoding is the genotypic representation that seems to have the best scalability properties among GA's. There are two types of grammar encoding: matrix grammar and graph grammar. Kitano proposed the matrix grammar encoding, in which the chromosomes encode a graph generation grammar [86]. A grammar rule rewrites a 1×1 matrix into a 2×2 matrix, and so on, until it generates a $2^k \times 2^k$ matrix, such that $2^k < N$, the maximum number of units in the network. An $N \times N$ matrix is extracted from the $2^k \times 2^k$ matrix. Each character is mapped into zero or one to generate the network's connectivity matrix. Unfortunately, this encoding is not very efficient since only a small percentage of the rules are used. Furthermore, in the case of feedforward NN's, less than half of the connectivity matrix is used. Recently, Siddiqi and Lucas [88] showed that grammar encoding is superior to direct encoding when the initial population of network topologies is formed with sparsely connected individuals (with a probability of connection = 0.3). However, the results produced by direct encoding for densely connected networks (with a

probability of connection = 0.7) were indistinguishable from Kitano's grammar encoding.

Grau proposed a graph grammar called cellular encoding [88]. He used an attribute grammar in which the grammar rules are represented by a tree structure. Each rewrite rule provides an operator that computes the attributes of a symbol on the rule's right-hand side from the attributes of the symbols on the rule's left-hand side. This approach was used to generate large-size Boolean networks.

2) *Using GP*: Koza and Rice proposed another solution to some of GA's perceived problems [89]. They advocated the use of GP, which does not rely on a predefined chromosome's length. GP uses a direct encoding of the NN's weights and topology in a genetic tree structure.

3) *Using EP*: The ability of determining the model's structure using evolutionary algorithms was proposed by Fogel [90] and McDonnell and Waagen [91], among others. The latter suggested using EP with a fitness function that combines a measure of the network's complexity (number of connections) with the mean sum-squared pattern error. This fitness favored removing those synapses that had small effect on the pattern error. Angeline *et al.* [59] argued that GA's were inefficient in evolving network topologies due to the deception problem [92]–[93] and the complexity of the interpretation function. Angeline proposed an EP, called generalized acquisition of recurrent links (GNARL). This system relies on a number of mutation operators that allow for the simultaneous structural and parametric optimization of the network. This evolutionary program was capable to generate networks with complex, nonsymmetrical topologies. Their view has been followed by many other researchers, such as Lee and Kim [94], who proposed the use of EP with a linked-list encoding to avoid the permutation problem caused by the many-to-one mapping between genotypes and phenotypes in GA's.

D. NN Controlled by FL

FL enables us to easily translate our qualitative knowledge about the problem to be solved, such as resource allocation strategies, performance evaluation, and performance control, into an executable rule set. This characteristic has been the basis for the successful development and deployment of fuzzy controllers.

Typically this knowledge is used to synthesize fuzzy controllers for dynamic systems [1]. However, in this case the knowledge is used to implement a smart algorithm controller that allocates the algorithm's resources to improve its convergence and performance. As a result, fuzzy rule bases and fuzzy algorithms have been used to monitor the performance of NN's or EC and modify their control parameters. For instance, FLC's have been used to control the learning rate of NN's to improve the crawling behavior typically exhibited by NN's as they are getting closer to the (local) minimum.

The learning rate is a function of the step size and determines how fast the algorithm will move along the error surface, following its gradient. Therefore the choice of the learning rate has an impact on the accuracy of the

final approximation and on the speed of convergence. The smaller its value the better the approximation but the slower the convergence. The momentum represents the fraction of the previous changes to the weight vector, which will still be used to compute the current change. As implied by its name, the momentum tends to maintain changes moving along the same direction thus preventing oscillations in shallow regions of the error surface and often resulting in faster convergence. Jacobs established a heuristic rule, known as the *delta-bar-delta* rule, to increase the size of the learning rate if the sign of the error gradient was the same over several consecutive steps [95]. Arabshahi *et al.* developed a simple fuzzy controller to modify the learning rate as a function of the error and its derivative, considerably improving Jacobs' heuristics [96].

The selection of these parameters involves a tradeoff: in general, large values of learning rate and momentum result in fast error convergence, but poor accuracy. On the contrary, small values lead to better accuracy but slow training, as proved by Wasserman [97].

E. EC Controlled by FL

The use of FL to translate and improve heuristic rules has also been applied to manage EC resources (population size, selection pressure, probabilities of crossover and mutation), during their transition from exploration (global search in the solution space) to exploitation (localized search) in the discovered regions of that space that appear to be promising [68], [70].

The management of EC resources gives the algorithm an adaptability that improves its efficiency and convergence speed. Herrera and Lozano suggest this adaptability can be used in the GA's parameter settings, the selection of genetic operators, solution representation, and fitness function [98].

The crucial aspect of this approach is to find the correct balance between the computational resources allocated to the metareasoning (e.g., the fuzzy controller) and to the object-level problem solving (e.g., the GA). This additional investment of resources will pay off if the controller is generic enough to be applicable to other object-level problem domains and if its run-time overhead is offset by the run-time performance improvement of the algorithm. An example of an application of this type of hybrid system can be found in [99].

F. Advantages of SC Hybrid Systems

This brief review of hybrid systems illustrates the interaction of knowledge and data in SC. To tune knowledge-derived models we first translate domain knowledge into an initial structure and parameters and then use global or local data search to tune the parameters. To control or limit search by using prior knowledge we first use global or local search to derive the models (structure + parameters), we embed knowledge in operators to improve global search, and we translate domain knowledge into a controller to manage the solution convergence and quality of the search algorithm.

All these facets have been exploited in some of the service and control applications described in the following sections.

III. SC APPLICATIONS FOR PREDICTIVE MODELING: DIAGNOSTICS, CONTROL, AND ESTIMATION

The task of predictive modeling is to forecast a course of events from a set of observations. For instance, predicting aircraft engine performance from exhaust gas temperatures (EGT) and fan speeds, or predicting the value of a residential property from the age and location of the house. In essence, predictive modeling synthesizes the nonlinear mapping from the inputs to the outputs. In the above examples, the input is the EGT and fan speed of an engine, or the age and location of a house, while the output is the engine performance or the value of the property, respectively.

Predictive modeling is the focus of many industrial and commercial applications. It can be used for diagnosis in which a set of fault hypotheses is identified—e.g., EGT exceeds its margin based on aircraft engine performance—and the system solves a classification problem (as in the first application that we discuss in Section IV). It can also be used for control in which the system has to output control actions that will impact the future state of the system (as in our second example, involving freight train control in Section V). Predictive modeling can also be used for estimation of the market value of a residential property, based on the attributes of the house (as in the last application, in Sections VI and VII).

IV. CLASSIFYING GAS TURBINE ANOMALIES WITH ADAPTIVE FUZZY CLUSTERING

A. Problem Description

Gas turbines are used in applications such as stationary power plants, airplanes, and helicopters. Depending on the particular use, the design and size of the gas turbine will vary. On a coarse level, however, gas turbines use the same operating principles. It is desirable to be able to detect incipient failures before they cause costly secondary damage or result in equally undesired shutdowns.

Service providers have long tracked the behavior of engines by measuring many system parameters and by using trend analysis to detect changes that might be indicative of developing failures [100]. Challenges of these schemes are that faults may not be recognized distinctively due to large amounts of noise as well as changing operating conditions that constantly move the estimate for “normal” operation. The former are caused in part by changing environmental conditions for which corrections with first principle models or regression models work only to some extent. The latter are due in part to changes of schedules, maintenance, etc., which are not necessarily known to the analyst.

In trend analysis, online sensor data and parameter settings are evaluated for deviations from normal operating conditions. When threshold settings are tripped and if the reading is not considered an outlier, an alarm is raised.

The settings used in trend analysis are typically below the limit that would cause an immediate shutdown and are meant to create awareness of potential problems before they turn into events that result in revenue loss (for airplanes, that could mean grounding of an airplane, damage to the engine, departure delays, etc.) An alarm normally causes the analyst to examine the trend charts to see if it appears that something real has happened. If something does appear to be suspicious, remedial actions (borescoping, engine washing, overhauls, etc.) are suggested. Problems with the system include the improper identification of alarms with engines where data quality is “bad” (data quality varies considerably between different engines due to different operating conditions and data acquisition) and where there is data drift due to wear. In addition, noise from poorly calibrated and deteriorating sensors and faulty data acquisition, among others, can result in excess generation of some alarms. If too many alarms are generated, the user must look at trends more often than desired to weed out the faulty alarms from the true ones. Increasing the thresholds typically will also increase the false negatives (missing an alarm condition) rate, which is typically even less desirable than a large number of false positives.

B. Approach

The approach for adaptive clustering combines multivariate data evaluation using fuzzy clusters with an exponential filter, which lets the clusters adapt to changing environments. The adaptive properties of the clusters allow the centroids to move and their shape to vary.

C. Prior Work

Karamouzis and Feyock proposed a system integrating model based and case-based reasoning techniques [101]. This system contained a self-organizing memory structured as a frame-based abstraction hierarchy for storing previously encountered problems. The cases were aircraft accident reports. Sensor inputs were gas temperature, engine pressure ratio, and human input such as sound of explosion and smell of smoke in passenger cabin. The time sequence was also of importance in establishing a causal relation such as an aberrant reading of the fan speed before an abnormal observation of the compressor speed. The system incorporated a functional dependency implemented via a digraph and a causality submodel describing transitions between various states of the system. The system searched its case library for the most similar case by a weighted count of corresponding symptoms and gave higher degree of similarity for symptoms occurring early in the fault occurrence.

Reibling and Bublin proposed a system that integrates pattern-matching techniques with causal networks to model normal and abnormal behavior [102].

Fernandez-Montesinos used Kohonen’s feature maps and expert systems to support the interpretation of in-flight monitoring data and to ensure more consistent engine condition monitoring (ECM) [103]. ECM was supported

by performance trend analysis software (ADEPT) which captured deterioration and failure of engine parts. Shift patterns allowed the localization of problems within the engine modules. The two main tasks tackled were recognition of a pattern indicating a possible problem, and the interpretation and further analysis. Kohonen's feature maps were chosen because of their self-organizing properties and the incremental extension of the domain of the patterns. Output of the Kohonen's feature map was an error code that indicated whether the module recognized a problem. As starting points for the network, fingerprints provided by the engine manufacturer (GE) were used. An expert system controlled the proposed neural net and interpreted the output.

Records and Choi [104] investigated strategies to filter spurious symptoms in aircraft engine fault monitoring systems (resulting from the monitoring system's inability to accurately assess the expected value from an engine model). They used a back-end knowledge-based approach and a front-end NN that generated expectation values for the monitored sensor. Gomm [105] and Patel [106] both suggested the use of a self-adaptive NN with online learning capabilities. Gomm used RBF networks to which new output nodes were automatically added when a new process fault is encountered. Adaptation was achieved using recursive linear algorithms that train the localized network parameters.

D. Solution Description

1) *Fuzzy Clusters*: Evaluation in multivariate feature space helps in distinguishing some faults because system variables are interconnected and changes in one variable may result in changes of other variables as well. For example, engine variables such as EGT, fuel flow (w_f), and turbine speed (n_2) are correlated such that particular failure conditions show up in changes of several variables. A bleed problem, where air leaks from the compressor, will result in a less efficient engine. However, in case of an aircraft engine, the controller demands a certain thrust level and will cause more fuel to be injected into the combustors. This in turn will raise the turbine speed as well as the exhaust gas temperature. The change in any of the engine parameters may be too small to be recognized alone. However, the change is more pronounced in a higher dimensional space and should therefore be detected more easily. Still, depending on the amount of noise, an overlap remains in between the cluster centers where variables cannot be partitioned clearly and the potential for misclassification remains high, making a crisp classifier not very feasible. The answer is not necessarily adding more features to the classification scheme because the decreasing potential information gain (in Euclidean space) is blurred additionally by decreasing feature quality, assuming one starts out using the highest-quality feature and adding others in decreasing order of value. At a certain threshold, more features may actually decrease the performance.

Engine behavior is clustered into different regions in the EGT- w_f - n_2 space. Data behaving normally will appear

in a cluster labeled "normal." The presence of an alarm condition, such as a rapid change of EGT caused by a compressor leak is located in the leak fault cluster. Rate of change of the data is not included here using the assumption that the abnormal condition will show up at least an order of magnitude faster than a slow drift (which is considered normal). The clusters are of nonuniform and nonlinear degrading size and shape. This allows data of a training set with a nonlinear boundary between classes to be separated more easily. Moreover, the boundaries between the clusters are not crisp. That is, the evaluation of the cluster membership is carried out such that the degree of membership is largest at the cluster center, with typically (but not necessarily) nonlinear slope outward. We use a fuzzy c -means clustering algorithm [107] for this purpose.

2) *Adaptation*: Normal behavior is constantly redefined. Thermal, chemical, and mechanical wear degrade the performance of the engine. These changes are expected but not necessarily predictable because they are driven partly by external factors. In addition, maintenance, controller software changes, and other factors may vary the operating conditions as well. At any rate, these changes must also be reflected in changes of the clusters to retain desirable classification properties. To achieve this goal, the centroid was superimposed with an exponential filter with small smoothing parameter, which forces the centroid to be updated according to the current data. Data were used for updating only when they were classified as part of a particular cluster. That is, data of type "normal" were used to update the cluster "normal" but not the cluster "leak fault" and vice versa. After updating of the memberships in the cluster (which calculates membership values and cluster centers), the adaptation of a fuzzy cluster is performed by borrowing the exponential weighted moving average filter from statistical signal processing, which we express as

$$\text{new } \check{v} = \text{old } \check{v} + \check{\alpha} \left(\frac{\sum_{i=1}^n \check{u}_i^m y_i}{\sum_{i=1}^n \check{u}_i^m} - \text{old } \check{v} \right)$$

where superscript "win" denotes the winner, superscript "lose" denotes the loser, $\text{new } \check{v}$ is the new cluster position, $\text{old } \check{v}$ is the old cluster position, $\check{\alpha}$ is the adaptive learning coefficient calculated (for two clusters) by $\check{\alpha} = \alpha(1 - |\check{u} - \hat{u}|)$, where α is a constant term, e.g., $\alpha = 0.05$, u is the membership value, and y_i is the data point under consideration

The adaptation rate, given by $\check{\alpha}$, tries to minimize ambiguity. When the degree of membership in a winning cluster is medium (ambiguous), \check{u} and \hat{u} are almost the same. In that case, $\check{\alpha}$ is large, thus trying to untangle the indeterminate situation, and moving the winning cluster farther toward the new data point. Moreover, it causes a higher rate of change near the true fault cluster boundaries and less change where the fault classification is already working well. Thus, it has the effect of focusing its learning where it is needed most, as dictated by recent data. As a




Fig. 2. Fuzzy clusters for several conditions displayed in EGT- n_2 space.

result, the centroids move with changing normal conditions. The scheme will dampen the effects of noise as well, since it gives some weight to past data and acts as an averaging scheme. It will also, like all filters, lag behind true changes in the system.

3) *Alertness Filter*: Another approach to increase robustness is to track the cumulative occurrence of “suspicious” readings, rather than tripping the alarm after just one “fault” classification. We propose an alertness filter for doing this, which forces the reoccurrence of several suspicious readings before an alarm is issued. If only one suspicious reading is encountered, the value of an “alertness counter”—which would be zero under normal conditions—is incremented. If a second suspicious reading is encountered immediately afterwards, the counter value goes up. Otherwise, it goes down. In the latter case, the first suspicious reading is treated just as an outlier and no maintenance is recommended. If it is not an outlier, i.e., the condition prevails, the level of the alertness will go up until a full alarm is issued, at which point remedial action will be advocated. At the alarm threshold, further fault readings will not increase the alertness level. Instead, it will be limited to a user-specified number. This saturation avoids lengthy recovery from alarm conditions when the system recovers either spontaneously or through maintenance. The purpose of the alertness filter is to further improve the classification rate of the system.

V. RESULTS

The classification system was tested with historical gas turbine data. Different events of similar fault conditions were examined and divided into training and test cases. For training, the use of both deliberate placement and

random seed points (for the cluster centroids) was investigated. Placing the starting points into the general expected neighborhood resulted in faster learning and convergence. Convergence was on average accomplished after about 20 iterations, depending on the start conditions. To prevent a fitting bias for normal conditions due to a vastly larger amount of data for that state, those data were undersampled so that “normal” and “faulty” operation were equally represented. Fig. 2 shows the distribution of clusters after training in EGT- n_2 space.

After training, the system was used for classification of online test data. The false negative rate is as low as with traditional trending tools (none observed with the test data available) and the false positive rate improved from 95% to less than 1%. Fig. 3 shows the adaptation of a cluster during operation.

The “ x ” denotes the location of the centroid of cluster “normal” and the squares denote the location of a fault cluster. During a software change, the normal region changes considerably in one dimension as seen by the vertical jump of the operating point upward in the graph. Because there were no changes in the other two dimensions, the algorithm adapts correctly to the new location while retaining the ability to detect faults. The drastic change also shows the lag introduced by the exponential filter.

A. Summarizing the SC Application to Diagnostics

We have presented the use of adaptive fuzzy clusters for fault classification of gas turbine engine sensor data. The advantages of this system are its ability to deal with extremely noisy sensors and to adjust to unpredictable slow drift. Dramatic improvement in performance was demonstrated during testing. Future work could focus on

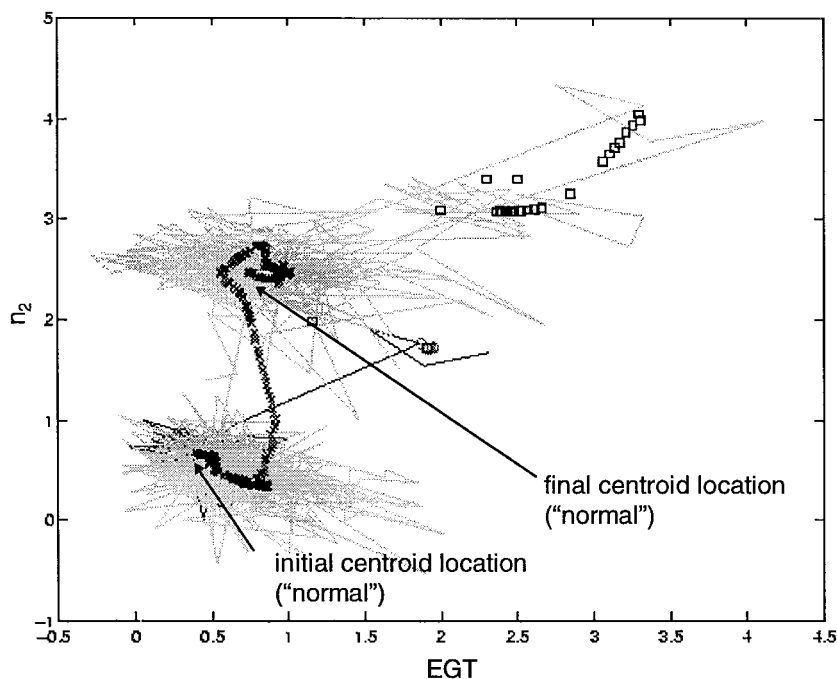


Fig. 3. Adaptation of clustering during operation.

detecting new fault types and the need for adaptively adding clusters.

In the broader context of SC and the theme of model = structure + parameters, the model structure here is provided by the fuzzy clusters (their number, which is a structural feature of the model). The model parameters are the centroid locations and membership functions, which are determined and tuned by a simplified Kalman filter type approach (exponential smoothing). In general, this type of search method is not provably optimal, unless certain assumptions about the system noise and behavior are met. The hybridization is a complementary one, in that two different technologies are achieving two functionally different aims.

Of course, it is possible to use other SC methods for this type of problem. For instance, EC could be used to determine the optimal centroid cluster locations in an offline computation using a limited set of initial data, but ongoing, real-time tuning of the clusters is currently impractical (if not unfeasible) if done by evolutionary algorithms. The clustering could be implemented as an RBF NN with competitive learning schemes (such as learning vector quantization and its variants). Probabilistic schemes (with subjective probabilities) could also be employed here for ranking cluster candidates, and indeed BBN's have been used widely in diagnostic applications in industry. The fuzziness used here lends itself to human interpretation of the clusters in some cases, which is an added benefit.

VI. AUTOMATED TRAIN HANDLING USING GENETICALLY TUNED FUZZY SYSTEMS

A. Problem Description

The development of an automated train handler requires the control of a massive, distributed system with little

sensor information. Freight trains consist of several hundred heavy railcars connected by couplers. A typical freight train can be as long as two miles and requires up to four coupled locomotives (12000 hp) to pull it. Each coupler between railcars may have a dead zone and a hydraulically damped spring. This implies that the railcars can move relative to each other while in motion, leading to a train that can change its length by 50–100 ft. The controlled forces on the train are due to a throttle and dynamic brake exerting a force on the locomotive, which is transmitted through the couplers, and a distributed air brake. Handling of these controls has a direct effect on the intercar coupler dynamics and the forces and distances therein (called slack). The position of the cars and couplers cannot be electronically sensed. Couplers are subjected to high static forces and dynamic forces, which may lead to breakage of the coupler, the brake pipe, and the train. Violation of speed limits and excessive acceleration/breaking may lead to derailment and severe cargo damage. Smooth handling while following a speed target is therefore absolutely imperative. These boundary conditions make it hard to control the train and the current completely manual control depends heavily on the experience of the crew. Current locomotives are equipped with a very simplistic cruise control that uses a linear proportional integral (PI) controller, which can be used only below speeds of 10 mi/h. This PI controller is primarily meant to be used for uniform loading, yard movement, etc., and does not prescribe braking action. Furthermore, the technology used does not consider slack or distributed dynamics in any way and is inappropriate for extended trains at cruising speeds over general terrain.

An automated system has to satisfy multiple goals: providing a certain degree of train handling uniformity

across all crews; enforcing railroad safety rules (such as posted speed limits); maintaining the train schedule within small tolerances; operating the train in fuel-efficient regimes; and maintaining a smooth ride by avoiding sudden accelerations or brake applications. This last constraint will minimize damage due to poor slack handling, bunching, and run in. As described above, the handling of freight trains involves a multibody problem and proper slack management, without sensors for most of the state. These constraints lead to a complex problem that cannot be solved by the simpler schemes used by cruise controllers for other vehicles, such as cars, trucks, boats, etc.

B. Approach

Our design uses a fuzzy controller, composed of a cruise planning module and a cruise control module that can automate the controls of a freight train [74]. This system would be applicable during most of the train journey, except for the initial and final transients, i.e., the train starting and stopping. The planning module focuses on creating a good reference trajectory that the fuzzy controller tries to track. Here, we focus only on the use of a GA to tune the fuzzy controller's performance by adjusting its parameters (the scaling factors and the membership functions) in a sequential order of significance. We show that this approach results in a controller that is superior to the manually designed one. To test and tune our fuzzy controller, we have used a simulator developed in-house, based on work done at GE and the Association of American Railroads.

C. Prior Work

There is no significant prior work on automatic train handling for the type of general terrain use that we are targeting here. Our proposed solution involves a hybridization of fuzzy control and genetic algorithms. Please refer to Section II-B for a review of the prior work in GA-based tuning of fuzzy controllers.

D. Solution Description

1) *Architecture*: The overall scheme for the proposed train handling is shown in Fig. 4. It consists of a fuzzy proportional integral controller (FPI) closing the loop around the train simulation (TSIM), using only the current velocity as its state input. This velocity and look ahead are compared with the desired profile to generate a predicted near-term error as input to the FPI, which outputs control actions back to the simulator. Offline, a GA uses the setup for evaluating various choices for the FPI parameters. The simulator TSIM is an in-house implementation, combining internal data with physical and empirical models.

The FPI controller uses the tracking error (e) and its error change (Δe) to recommend a change in the control outputs throttle notch and brake settings (Δu). If allowed, the FPI will track the reference profile accurately. The computation of the error used by the FPI incorporates

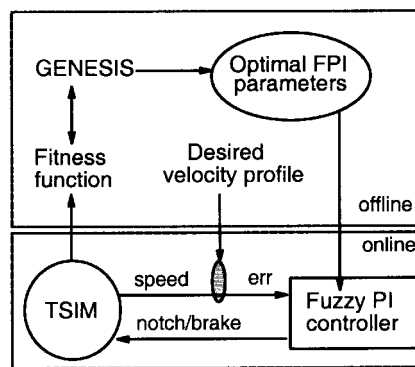


Fig. 4. System schematic for using a GA to tune a fuzzy train controller.

a look-ahead to properly account for the train inertia. The algorithm predicts the future velocity of the train and incorporates not only the current error, but also the future predicted error, since the future reference velocity is known from the profile. Finally, the PI is tuned (offline) using GA's that modify the controller's most sensitive parameters: SF's and membership function parameters (MF).

2) *Designing the Fuzzy PI Controller*: To summarize the controller briefly, it is described by a KB composed of an RS, TS's of MF's, and SF's. The RS maps linguistic descriptions of state vectors [$e; \Delta e$] into incremental control actions Δu ; the TS's define the semantics of the linguistic values used in the rules; and the SF's determine the extremes of the numerical range of values for both the input and output variables. The relationship between output variable, u , and control error, e , can be expressed approximately as follows [75]:

$$\begin{aligned} \frac{\Delta u(t)}{S_u} &\approx \frac{\Delta e(t)}{S_d} + \frac{e(t)}{S_e} \\ u(t) &\approx \frac{S_u}{S_d} e(t) + \frac{S_u}{S_e} \int e(t) dt \\ -S_e &\leq e(t) \leq S_e \\ -S_d &\leq \Delta e(t) \leq S_d \\ -S_u &\leq \Delta u(t) \leq S_u \end{aligned}$$

where S_e , S_d , and S_u are the scaling factors of error, change of error, and incremental output variable, respectively. On the other hand, a conventional PI controller has $u(t) = K_p e(t) + K_i \int e(t) dt$, where K_p and K_i are the proportional and integral gain factors, respectively. Comparing the above equations, we obtain

$$K_p \approx S_u/S_d, \quad K_i \approx S_u/S_e.$$

3) *Tuning the Fuzzy PI*: For the purposes of this study, GENESIS (Genetic Search Implementation System) has been adopted as a software development tool [108]. The user needs to provide only a fitness function that returns a corresponding value for any point in the search space.

a) *Fitness functions:* To study the effects of different objectives, we investigated two fitness functions:

$$f_1 = - \left(\sum_i |\text{notch}_i - \text{notch}_{i-1}| + |\text{brake}_i - \text{brake}_{i-1}| \right)$$

$$f_2 = - \left(w_1 \frac{\sum_i |\text{notch}_i - \text{notch}_{i-1}|}{K_1} + w_2 \frac{\sum_i |v_i - v_i^d|}{K_2} \right)$$

where v^d denotes the desired reference velocity and i is a distance or milepost index. The function f_1 captures jockeying of the throttle and brakes. The function f_2 combines a weighted sum of velocity profile tracking accuracy and throttle jockeying. The train simulator TSIM provided the environment in which the train's response to each controller could be tested, so that the values of the fitness function could be calculated over the entire journey.

b) *Design choices:* A critical constraint in our design formulation was the cost of running the TSIM simulator. To analyze the behavior of the train dynamics properly, we needed to use a very small integration step. As a result, the accurate simulation of a 40-mi track required more than 12–15 s on a Sun Sparc10. Since a complete simulation was required during each trial, we wanted to limit the number of trials, which is the product of population size and number of generations, while still achieving a reasonable convergence. To improve the convergence, we decided to limit the chromosome length, which is one of the factors that affect convergence. Therefore, we opted for a sequential tuning approach of SF's, MF's, and RS's. Furthermore, we decided to use a coarse precision for each of the SF by using eight levels (three bits), since that granularity was enough to ensure the FLC's good performance.

There are other promising alternatives that we omitted in this first phase of the project, such as the simultaneous tuning of both SF's and MF's or the use of a finer granularity obtained by increasing the chromosome length. These alternative design choices will be addressed in our future work.

In this experiment, we performed eight tests by taking a cross product of the scaling factor values before and after GA tuning, the MF parameter values before and after GA tuning, and the two fitness functions.

c) *Train simulator parameters:* All testing for the automated tuning of FPI was done using TSIM. Two track profiles were used: an approximately 14-mi flat track (which captures travel in the plains) and an approximately 40-mi piece of actual Conrail track from Selkirk to Framingham over the Berkshires in Massachusetts (which captures mountainous terrain). The train weighed nearly 9000 tons, and was about 1 mi long, with four locomotives and 100 fully loaded railcars. An analytically computed velocity

profile that minimizes fuel consumption was used as the reference.

d) *FPI controller parameters:* The standard TS used in the FPI is {NH, NM, NL, ZE, PL, PM, PH} where N = negative, P = positive, H = high, M = medium, L = low, and ZE = zero. Initially, the terms were uniformly positioned trapezoids overlapping at a 50% level over the normalized universe of discourse. Since the controller was defined by a nonlinear control surface in $(e, \Delta e, \Delta u)$ space, we needed three TS's in all, one for each of these three variables. This design leads to a symmetric controller, which is not always a good assumption. In this case, the GA tuning automatically created the required asymmetry.

e) *GENESIS parameters:* The GA parameters were initialized according to the default values suggested by De Jong [109], i.e., population size = 50, crossover rate = 0.6, and mutation rate = 0.001. In addition, all structures in each generation were evaluated, the elitist strategy was used to guarantee asymptotic convergence [110], gray codes were used in the encoding, and selection was rank based.

f) *Tuning of SF's:* Each chromosome is represented as a concatenation of three 3-bit values for the three floating-point values for the three FPI scaling factors $S_e, S_d,$ and S_u . They are in the ranges $S_e \in [1, 9], S_d \in [0.1, 0.9], S_u \in [1000, 9000]$. The intent is to demonstrate a GA's ability for function optimization even with such a simple 9-bit chromosome.

g) *Tuning of MF's:* When tuning MF's, a chromosome is formed by concatenating the 21 parameterized MF's. Since each MF is trapezoidal and the overlap degree of 0.5 is maintained between adjacent trapezoids, this partitions the universe of discourse into intervals that alternate between being cores of a MF, and overlap areas. The core of NH and PH extends semi-infinitely to the left and right, respectively (outside the $[-1, 1]$ interval). These intervals, denoted by b_i are 11 in number for 7 MF labels. Typically, $\#(b) = 2 \times \#(\text{MF}) - 3$, where $\#(b)$ is the number of required intervals and $\#(\text{MF})$ is the number of membership functions. Each chromosome is thus a vector of 11 floating point values. Since the universe is normalized, $\sum_{i=1}^{11} b_i \leq 2$. For the present study, each b_i is set within the range of [0.09; 0.18] and five bits are used to represent a chromosome for GA-tuned membership functions. If $\sum_{i=1}^{11} b_i$ exceeds two, this implicitly reduces the number of effective MF's, thus providing partial structure optimization, as well.

E. Results

1) *Tuning SF Versus MF with Fitness Function f_1 :* Four sets of tests were conducted for the comparison of significance in tuning SF's versus tuning MF's with respect to fitness function f_1 . The results are shown in Table 1. The initial set of scaling factors is $[S_e S_d S_u] = [5.0; 0.2; 5000]$. After four generations of evolution, the set of GA-tuned SF is [9.0; 0.1; 1000]. As shown in Table 1, the fitness value was dramatically increased from -73.20 to -15.15 for the GA tuned SF with respect to f_1 . Substantially smoother control is the result.

Table 1
Summary of Simulation Results on 14-mi Flat Track

Description	f1	f2
Initial SF w/initial MF	- 73.20	- 1.00
GA tuned SF w/initial MF	- 15.15	- 0.82
Initial SF w/GA tuned MF	- 70.93	- 0.94
GA tuned SF w/GA tuned MF	- 14.64	- 0.82

On the other hand, GA-tuned MF only reduced throttle jockeying slightly after 20 generations of evolution. This is reflected in the small increase in fitness value from -73.20 to -70.93 . From the above observations, we can conclude that tuning scaling factors is more cost effective and has a larger impact than tuning membership functions. The small improvement by tuning MF's leads to an asymmetric shift in the membership functions, such that the ones corresponding to the brake lever are shifted more than the ones corresponding to the throttle setting. This is due to the slightly asymmetric calibration of the notch and brake actuators. A $+5$ (notch) may not lead to the exact tractive force forward that a -5 (brake) leads to in the opposite direction. As a result, the best FPI learns to respond to negative errors slightly differently than to positive errors.

Next, we demonstrate that tuning membership functions only gives marginal improvements for a fuzzy PI controller with tuned scaling factors. We used GA's to optimize FPI's MF's with respect to f_1 , while using the GA-tuned SF values of $[9.0, 0.1, 1000]$. After ten generations of evolution, the fitness value was further increased from -15.15 to -14.64 . The change in smoothness and the MF values is minimal. After observing the simulation results, we concluded that tuning membership functions alone provided only marginal improvements for a fuzzy PI controller with tuned scaling factors.

2) *Tuning SF Versus MF with Fitness Function f_2* : We further verified our arguments stated in Section V-E1 with the following four sets of tests conducted with respect to fitness function f_2 , which balances both tracking error and control smoothness.

Recall that the initial set of scaling factors is $[5.0, 0.2, 5000]$. After four generations of evolution, the GA discovered a set of SF = $[3.3, 0.9, 5571]$, with respect to function f_2 . As shown in Table 1, the fitness value was increased from -1.00 to -0.82 while doing this.

On the other hand, GA tuned MF only increased the fitness value from -1.00 to -0.94 after 20 generations of evolution, confirming the claim that tuning SF is more cost effective than tuning MF. We proceeded to experiment with MF tuning with tuned SF = $[3.3, 0.9, 5571]$. This time there were no significant improvements in fitness after ten generations.

Table 1 summarizes all eight tests run on the 14-mi flat track. In addition, we present the final performance graphs in Fig. 5 for the more complex piece of 40-mi real track with the same train. It shows substantial improvement in control accuracy and smoothness. Fig. 5(a) shows reference tracking performance and control outputs before GA tuning.

Fig. 5(b) shows vastly improved tracking and smoothness of control.

F. Summarizing the SC Application to Control

We have presented an approach that uses GA's to tune fuzzy systems for a complex control problem. We showed that all parameters do not need to be treated as equally important, and that sequential optimization, using GA's, can greatly reduce computational effort by tuning scaling factors first. Additional improvement was shown by the good performance of fairly coarse encoding. The scalability of the approach enables us to customize the controller differently for each track profile, though it does not need to be changed for different train configurations. In this way, we can efficiently produce customized controllers offline. Future work will also focus on automatic generation of velocity profiles for the train simulator by using GA's for trajectory optimization subject to "soft" constraints.

The problem context that we discussed resembles reinforcement learning, since the fitness function is cumulative over many output responses of the controller. The model uses fuzzy rules, but the rule structure is different from the fuzzy clusters used in the previous diagnostics problem. These rules are much more structured, with the standard three-dimensional FPI control surface and a large number of degrees of freedom due to the MF's. The parameter search space is also more complex, but the search method is more sophisticated. A simple filtering technique like the one used before would not work here because of the reinforcement learning issue and because the filter introduces a lag whereas this problem needs a look-ahead capability. On the other hand, the sequential optimization makes the search practical without sacrificing performance. Finally, the performance requirement is crucial enough that global search methods such as evolutionary algorithms are a better fit than the alternative gradient-descent based techniques. The hybridization is again of a complementary kind, where the EC algorithms can exploit the hierarchy of importance of the parameters in the FPI, unlike an NN where it is not clear if any weights could be tuned prior to the others.

VII. VALUING RESIDENTIAL PROPERTIES USING A FUZZY-NEURAL SYSTEM

A. Problem Description

Residential property valuation [111] is the process of determining a dollar estimate of the property value for given market conditions. For this paper we restrict ourselves to a single-family residence designed or intended for owner occupancy. The value of a property changes with market conditions, so any estimate of its value must be periodically updated to reflect those market changes. Any valuation must also be supported by current evidence of market conditions, e.g., recent real estate transactions.

The current manual process for valuing properties usually requires an on-site visit by a human appraiser, takes several

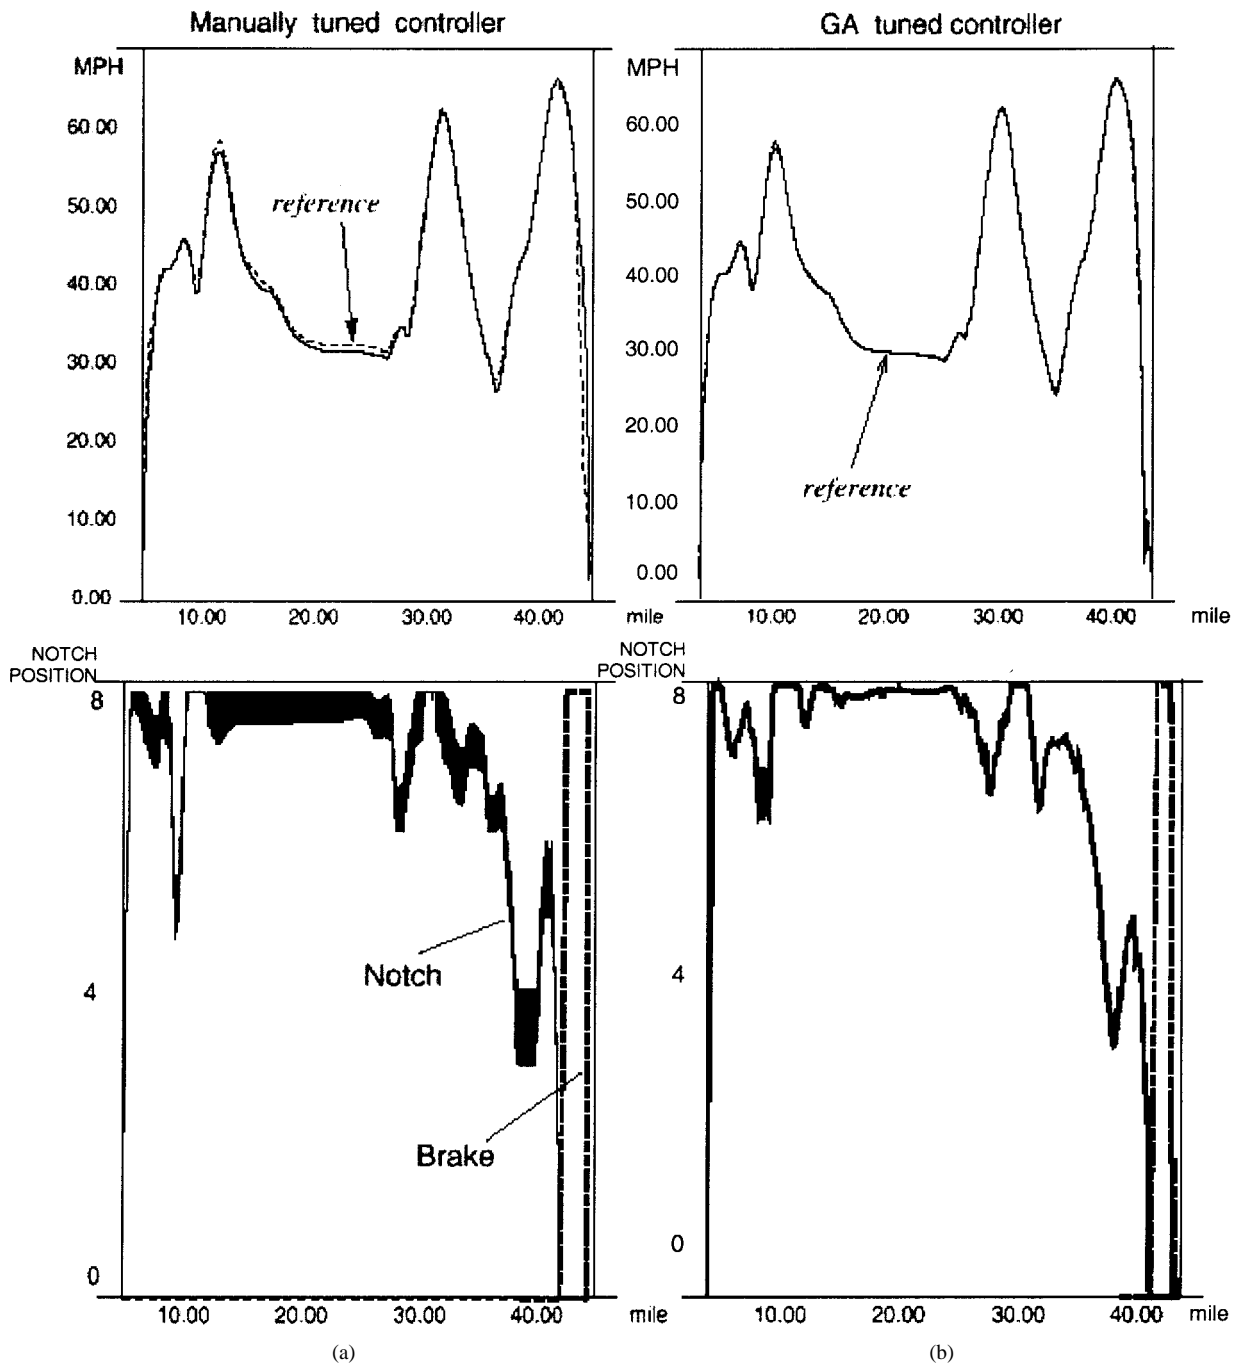


Fig. 5. Performance graph with manually tuned and GA tuned.

days, and costs about \$500 per subject property. This process is too slow and expensive for batch applications such as those used by banks for updating their loan and insurance portfolios, verifying risk profiles of servicing rights, or evaluating default risks for securitized packages of mortgages. The appraisal process for these batch applications is currently estimated, to a lesser degree of accuracy, by sampling techniques. Verification of property value on individual transactions may also be required by secondary buyers and mortgage insurers. Thus, this work is motivated by a broad spectrum of application areas. Some of the applications also require that the output be

qualified by a reliability measure and some justification, so that questionable data and unusual circumstances can be flagged for the human who uses the output of the system.

The most common and credible method used by appraisers is the sales comparison approach. This method consists of: finding comparables, i.e., recent sales that are comparable to the subject property (using sales records); contrasting the subject property with the comparables; adjusting the comparables' sales price to reflect their differences from the subject property (using heuristics and personal experience); and reconciling the comparables adjusted sales prices to derive an estimate for the subject property (using any

Table 2
Subset of Property and Sale Attributes

Attribute	Value
Sale Price	\$ 175,500
Sale Date	7/23/93
Address1	12 Bronco Lane
City	Contra Costa
State	CA
Zip Code	94063
Living Area	2,200 sq. ft.
Lot Size	10,000 sq. ft.
Bedrooms	3
Bathrooms	2
Total Rooms	6
Age	10 years

reasonable averaging method). This process assumes that the item's market value can be derived by the prices demanded by similar items in the same market. To automate the valuation process we have developed the Automated Residential Property Valuation (ARPV) system, which combines the result of two independent estimators—one using neurofuzzy networks, and one using case-based reasoning [112].

First, we will discuss the data used to develop and test our algorithm, followed by a brief overview of the methods employed to build individual estimators. We will then describe in detail a neurofuzzy-based estimator. Finally, we will present a method for fusing the output of the estimators into a single estimate and reliability value.

1) *Data Description:* The database used consists of public sales records of every property that has been sold in certain counties in Northern and Southern California during the five years preceding the development work. Data were purchased from two sources, which we shall name A and B. Data vendors provided a description of each property, which, in some cases, contained up to 166 property attributes. In addition, sales information (such as price and date) was also provided. A subset of the attributes describing a property is given in Table 2.

B. AIGEN Estimator

The generative artificial intelligence (AI) model (AIGEN) relies on a fuzzy NN that, after a training phase, provides an estimate of the subject's value.

The specific model that we chose for the AIGEN estimator is a network-based implementation of fuzzy inference. Part of it is based on ANFIS [67], which implements a fuzzy system as a five-layer NN so that the structure of the net can be interpreted in terms of high-level rules. This net is then trained automatically from data. We developed an extension of this methodology (referred to as E-ANFIS) which we use in AIGEN. ANFIS and E-ANFIS implement the fuzzy reasoning process as an NN and rely minimally on expert knowledge and mostly on market data. In essence, the FL helps to specify the structural component of the model only.

1) *E-ANFIS:* Fig. 6 shows the architecture that we used for the automated property valuation problem. The central difference between ANFIS and E-ANFIS is that E-ANFIS allows the output to be linear functions of variables that do not necessarily occur in the input. Another way of stating that is to allow the segmentation of the input space on a proper subset of the total variable set only and then using a cylindrical projection of that segmentation for the whole space.

Fig. 7 shows a schematic for the fuzzy inference process where the rules are TSK-type. The semantics are of the form: "if *lot_size* is small and *living_area* is small and *locational_value* is high, then price is $f(\dots)$," where $f(\dots)$ is a linear function of the 6 input variables. A special case of this is when all c_{ij} except c_{i0} are zero, in which case each rule recommends a fixed, crisp number. The inference procedure with TSK-type rules yields

$$z = \sum_i w_i (c_{i0} + c_{i1}x + c_{i2}y) / \sum_i w_i$$

where w_i is the weight of rule i , computed as a weighted sum. For the antecedent fuzzy membership functions, we use a generalized bell function (softened trapezoids) given by

$$\mu_A(x) = 1 / (1 + ((x - c)/a)^{2b}).$$

The E-ANFIS network architecture is determined by the number of membership functions assigned to each input dimension. For instance, if there are six inputs in all, and two membership functions are assigned to each of four of the inputs, the net has six input units, eight units in the first layer (which come from the two fuzzy functions for each of the four variables), 16 in the next two layers (since $2 \times 2 \times 2 \times 2 = 16$ rules), and finally, one summation unit to produce the output in the output layer. Each of the 16 rules has a TSK-type consequent that depends on all six inputs. Since each antecedent membership function has three degrees of freedom (a , b , and c), and each consequent has 7 coefficients here, there are $8 \times 3 + 16 \times 7 = 136$ degrees of freedom for this E-ANFIS model.

Once the architecture is constructed in this fashion, the parameters can be initialized with reasonable values, representing current domain knowledge, instead of randomly as in NN's. For instance, the membership functions can be spaced at uniform distances over the axis so as to cover the range of the data points. The consequent linear functions are initialized to zero.

2) *Training E-ANFIS:* A variant of the gradient descent technique is used to train the network based on the training data. In brief, this algorithm tries to minimize the mean squared error between the network outputs and the desired answers, when presented with the data points in the training set. It proceeds as follows.

- 1) Present a sample point in the training data set to the network and compute its output.
- 2) Compute the error between the network output and the desired answer.

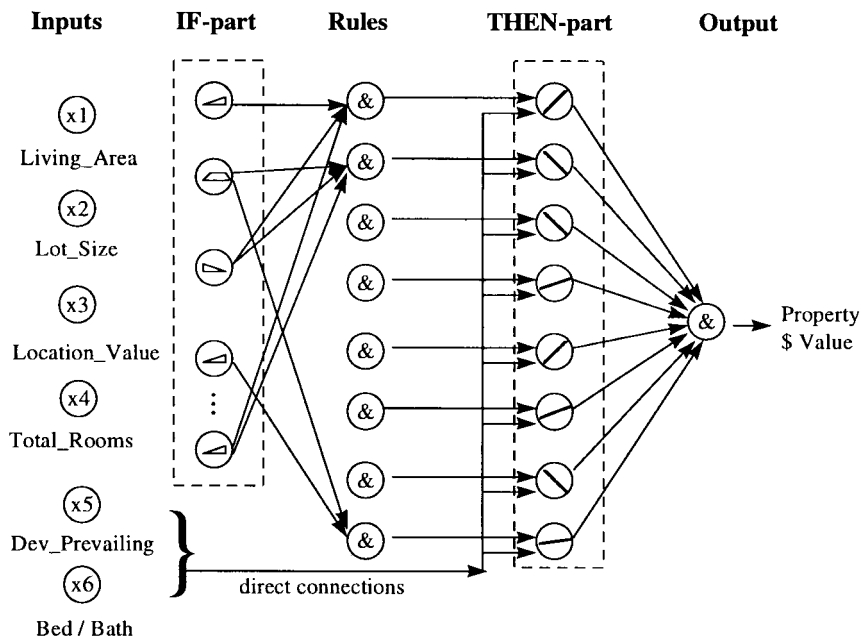


Fig. 6. An E-ANFIS network.

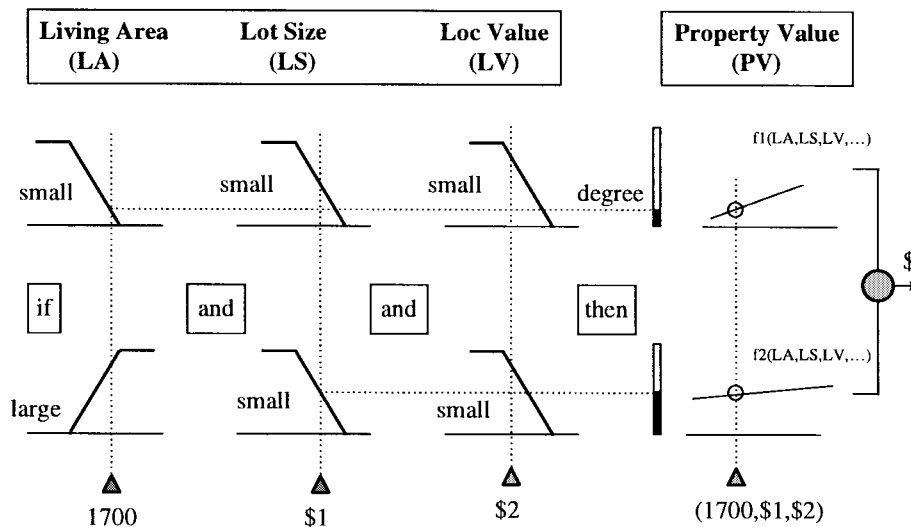


Fig. 7. The fuzzy inference process.

- 3) Holding the IF-part parameters fixed, solve for the optimal values of the THEN-part parameters using a least-mean-squares optimization method. A recursive Kalman filter method is used here.
- 4) Compute the effect of the IF-part parameters on the error, using derivatives of the functions implemented by intermediate layers.
- 5) Using the above, change the IF-part parameters by small amounts so that the error at the output is reduced.
- 6) Repeat the above steps several times using the entire training set, until the error is sufficiently small.

This procedure is as discussed in [67], with the extension of the Kalman filter to input variables that are not part of the antecedents of the fuzzy rules.

This procedure converges to a locally optimal configuration where the error becomes fixed or decreases very

slowly. The resulting network can be interpreted as a fuzzy rulebase, with each parameter in the net having a definite meaning in terms of the fuzzy sets or consequent functions. Learning speed is very fast compared to the conventional NN paradigm. Additional data, if available, can always be used to further train the model using the same BP-type algorithm. The resulting surface is well behaved and provably smooth. The rule base is extremely compact, so a large number of models can be stored easily. This is an important consideration, since there are thousands of counties in the United States, and each one will need at least one model.

The properties used for training the AIGEN model were restricted to be in a certain price range. This is to eliminate obvious outliers. For this reason, it should not be used to make a prediction on a property outside this range. It will make a prediction if given such a property, but will bound

its output to the said range. It will also issue a warning to the user in this case.

3) *Constructing AIGEN Models*: The training algorithm used was the BP-like method described earlier. We reduced the computational cost associated with training AIGEN by using a heuristic to reduce the dimensionality of the parameter space during training.

The 16 rules have 112 degrees of freedom in the consequent. This is a large share of the dimensionality of the parameter space, which uses a variant of the Kalman filtering algorithm to train the parameters in the consequent. It was found easier to train the consequent partially (4 inputs = 80 parameters) in the interleaved BP process described earlier, followed by a final batch phase where all 112 consequent parameters were retrained again while holding the antecedent parameters constant. This reduces the computational cost somewhat. This is similar in principle to the sequential optimization idea used for the train control in the previous application.

Recall that the model search process minimizes the mean squared error (MSE) over a training set at the present time. All the training, models, and results discussed here are with respect to this optimization metric. The choice of optimization criteria is pivotal to the model that is extracted by the process.

a) *Training set size*: We used about 5% of the total available data for training purposes. For one of our counties, this meant that out of the 37712 potential test records that were available from Data Source A, 1768 records were extracted and used for training the AIGEN model. For another data source, less than half of this (14–15 thousand records) was available, so that 10% of the total data was used for training the models. This corresponds to about 1400–1500 records for training.

These had to be error free and more or less randomly distributed so as not to bias the estimator. The specific size of the training set is not significant. Since there are 136 degrees of freedom in all, a good rule of thumb is to use at least ten times as many examples for training. Moreover, it is important to use a small part of the total data, so that overfitting to the data can be avoided. The results discussed below apply to this training scheme.

Inputs to the E-ANFIS network are based on seven attributes of the property:

- 1) *total_rooms*;
- 2) *num_bedrooms*;
- 3) *num_baths*;
- 4) *living_area*;
- 5) *lot_size*;
- 6) *locational_value*;
- 7) *deviation_from_prevaling_value*.

The number of bedrooms and bathrooms is combined to produce a *bedrooms/bathrooms ratio* that is fed along with the other five values to the E-ANFIS net. The last two variables are outputs from another estimator, LOCVAL, which is briefly described in Section VII-A1. These two variables serve as a rough starting point for the AIGEN

estimator. Of these six inputs, *total_rooms*, *living_area*, *lot_size*, and *locational_value* are used for partitioning the space into 16 fuzzy regions. The output variable is simply the dollar value of the house.

C. Results and Validation

Validation was done by testing the model on the entire dataset available from a data source, after it was filtered to remove atypical properties. The filters used for testing were the same as the ones used for screening the training set. Recall that the training metric was mean squared dollar difference between the actual and estimated price. We used the median of absolute relative error as the principal test metric, defined as the median value of $E = (\text{actual price} - \text{estimate})/\text{actual price}$. The median was chosen for its robustness to outliers.

The test dataset used for validation consisted of 35 370 property records from data source *A*. The training data size was 1768 points—only 5% of the total data.

The median value of E was in the range of [8.5%, 8.64%]. The average bias of the estimator was about \$557, which implies that the estimator was very slightly biased toward the conservative side, not an undesirable feature in this context. The bias shown by the unsigned relative error was higher, but this was purely because the same dollar value error translates into different relative errors, depending on the value of the property. The average E value was around 12%, since there are some extreme outliers in the test set.

Note that some limited amount of outlier elimination is took place in the model itself, since it bounded its own estimate to lie in the range used for training. Without the truncation, the median and mean of E were not significantly affected, and the root mean squared error (RMSE) increased slightly. This was to be expected, since it attached very high weights to the few far outliers. For test data in the [5%, 95%] percentile range of signed relative error, the coefficient of correlation between the actual price and the estimated AIGEN price was 0.97. Additional quality control of the data using some filters (on both the training and test sets), led to an improved median E of about 7–8%.

When data source *B* was used (which is of worse quality and completeness), the test set size was 15 733 records for the same county and time period. Training was done using 10% of these. The model performance with the same structure and search method was slightly worse, which is not unexpected.

We tried several alternative model structures as well, by varying the input variables fed to the net. The cumulative probability distributions for error E over the test sets (35370 points for *A*, 15733 for *B*) were compared. We confirmed that models based on dataset *B* were worse than the ones based on dataset *A*, but that the spread between model performances within each of the two sets was not very large (with *A* again producing more consistent models than *B*). Changing filters and price ranges improved the performance with respect to data set *B* as well.

We also did cross-checking experiments where a model trained on data source A was tested on data source B and vice versa. As expected, test performance deteriorated, partly due to data quality and partly due to inconsistent semantics used by the collectors of the two data source companies. This demonstrated the intuitive fact that when the model is used, the inputs to it should not only be accurate and complete, but consistent with the semantics of the training data.

D. Discussion

We used a complementary hybrid—a neurofuzzy system, to solve the property valuation problem in one particular way with AIGEN. Encouraging results were demonstrated in one California county, based on multiple data sources. This was despite the limited number of attributes used in the computation. These models were built for a few other counties as well, with comparable test results.

Considering that AIGEN looked only at five structural variables, and did not consider style, condition, quality of construction, and economic trend metrics, these were encouraging results. The other physical variables were not used since they tended to be missing from the records in a significant portion of the database.

Non-AI approaches such as statistical regression use a similar philosophy of training and testing the models. AIGEN can be thought of as a fuzzy extension of the CART tool [113], which uses an explicit tree structure as the model and the branching thresholds as parameters. The parameters are identified using a greedy variance-minimizing search algorithm. Our approach, in addition to being a smooth combination of piecewise linear regions, provides a transparent explanation of the model in terms of a small number of rules that are comprehensible to a human.

The hybridization in AIGEN is also complementary, where neural techniques help to define a fuzzy structure. As discussed earlier, small variations in model structure and inputs did not seem to affect the performance much. The number of nominal parameters was in the range of 100–150, but the number of effective parameters is probably less than that. The search method was entirely gradient based and could conceptually get trapped in a local minimum. However, there is independent evidence (from other estimators) that since this is a real-world illiquid market with noisy prices, there will always be an inherent irreducible error of about 5%, which a global search could not beat. Hence, using a local search technique here was much more practical than using genetic or evolutionary algorithms for search.

VIII. RESIDENTIAL PROPERTY VALUATION USING A FUSION OF HYBRID SOFT COMPUTING MODELS

In Section VI we focused on using a complementary hybrid of two SC technologies to solve a particular problem. If two or more such models could be built independently for the same problem, which one should be chosen? We propose that a loose hybrid, which could be called information fusion or model fusion, is a much more robust alterna-

tive to picking one of the models. We demonstrate this with the same application of residential property valuation where multiple estimators, based on location (LOCVAL), comparable properties (AICOMP), and a generative model (AIGEN), were built using different technologies on the same database. This combines the results of the models rather than the techniques themselves, and the focus is on using the redundancy to increase reliability rather than accuracy alone.

A. Estimators Used in Fusion

1) *Locational Value Model (LOCVAL)*: The first model was based solely on the location and living area of the properties, as shown in Fig. 8.

A dollar per square foot measure was constructed for each point in the county, by suitably averaging the observed, filtered historical market values in the vicinity of that point. This locational value estimator LOCVAL produces two output values: *Locational_Value* (a dollar/ft² estimate) and *Deviation_from_prevaling_value*. The local averaging is done by an exponentially decreasing radial basis function with a “space constant” of 0.15–0.2 mi. It can be described as the weighted sum of radial basis functions (all of the same width), each situated at the site of a sale within the past one year and having an amplitude equal to the sales price. Deviation from prevailing value is the standard deviation for houses within the area covered and is derived using a similar approach.

In order to use the LOCVAL estimator correctly, the input values (a valid, geocoded address and a living area in squared feet) must be present and accurate for the locational estimator to work correctly. If either is missing, or clearly out-of-range, the estimator will not make any prediction.

2) *Comparable Value Model (AICOMP)*: The second model, AICOMP, relied on a case-based reasoning (CBR) process similar to the sales comparison approach [111] used by certified appraisers to estimate a residential property’s value. The CBR process consists of selecting relevant cases (which would be nearby house sales), adapting them, and aggregating those adapted cases into a single estimate of the property value.

Gonzalez first showed the possibility of using CBR to automating the valuation process [114]. However, his CBR approach never captured the intrinsic imprecision of the basic steps in the sale comparison approach: finding the *most similar* houses, located *close* to the subject property, sold *not too long* ago; and selecting a *balanced* subset of the *most promising* comparables to derive the final estimate. Therefore we developed AICOMP, a fuzzy CBR system that uses fuzzy predicates and fuzzy-logic based similarity measures [115] to estimate the value of residential property. Our approach, shown in Fig. 9 and further described in [3], consists of the following.

- 1) *Retrieving recent sales from a case base*: Upon entering the subject property attributes, AICOMP retrieves potentially similar comparables from the case-base. This initial selection uses six attributes: address; date

Done before use

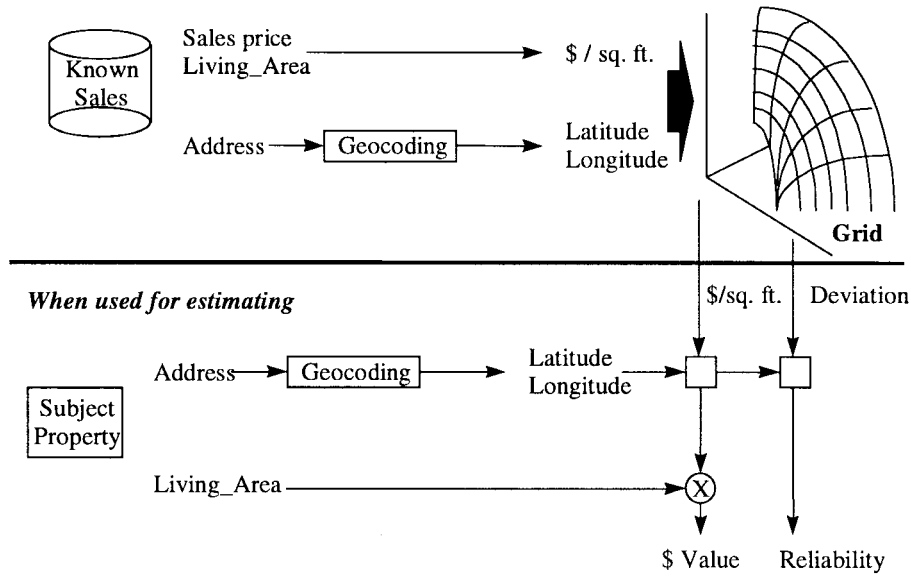


Fig. 8. Locational value method (LOCVAL).

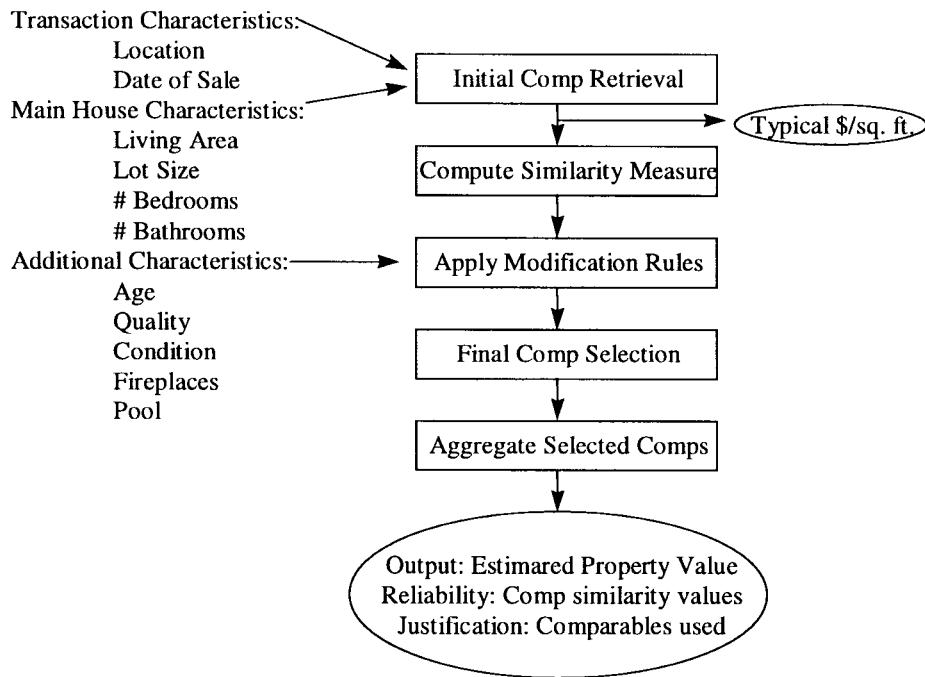


Fig. 9. CBR process.

of sale; living area; lot area; number of bathrooms; and bedrooms.

- 2) *Comparing the subject property with the retrieved cases:* The comparables are rated and ranked on a similarity scale to identify the most similar ones to the subject property. This rating is obtained from a weighted aggregation of the decision maker preferences, expressed as fuzzy membership distributions and relations.
- 3) *Adjusting the sales price of the retrieved cases:* Each property's sales price is adjusted to reflect their differences from the subject property. These adjustments

are performed by a rule set that uses additional property attributes, such as construction quality, conditions, pools, fireplaces, etc.

- 4) *Aggregating the adjusted sales prices of the retrieved cases:* The best four to eight comparables are selected. The adjusted sales price and similarity of the selected properties are combined to produce an estimate of the subject value with an associated reliability value.

3) *Generative Method (AIGEN):* The third method, called AIGEN, is a generative AI method in which a fuzzy

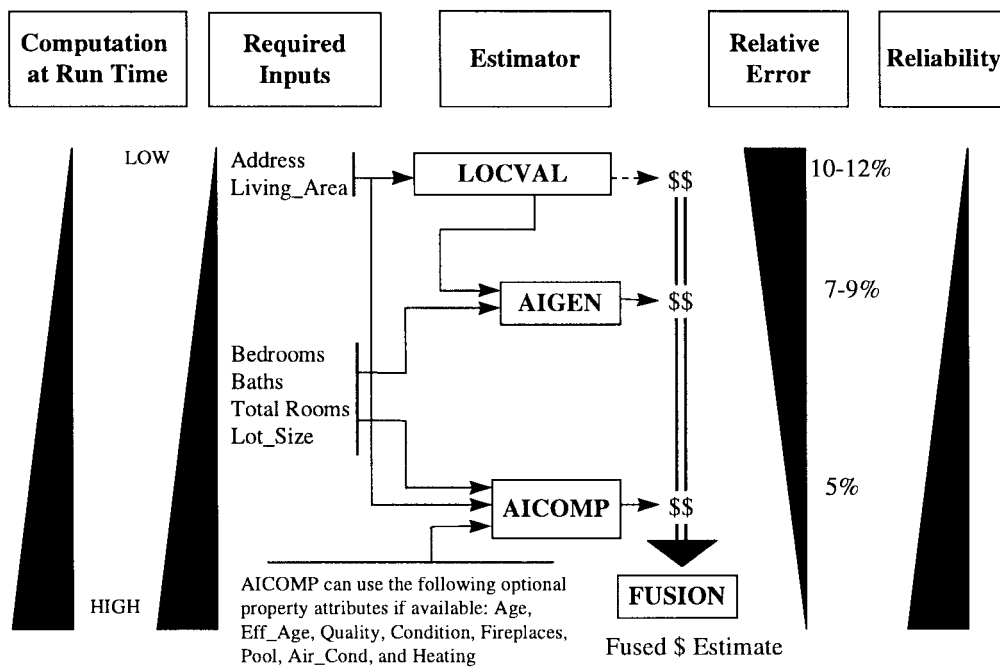


Fig. 10. Data comparison of multiple approaches.

NN is trained by using a subset of cases from the case base. The resulting run-time system provides an estimate of the subject's value. This method was described in Section VI.

4) *Model Outputs—Property and Reliability Values:* Each model produced a property value and an associated reliability value. The latter was a function of the “averageness” or “typicality” of the subject property based on its physical characteristics (such as lot size, living area, total room). These typical values were represented by possibilistic distributions (fuzzy sets). We computed the degree to which each property satisfied each criterion. The overall property value reliability was obtained by considering the conjunction of these constraint satisfactions (i.e., the minimum of the individual reliability values). A more detailed description of this process can be found in [116].

The computation times, required inputs, errors and reliability values for these three methods are shown in Fig. 10. The locational value method takes the least time and information, but produces the largest error. The CBR approach takes the largest time and number of inputs, but produces the lowest error.

5) *Advantages of Fusion:* The fusion of the three estimators has several advantages.

- 1) The fusion process provides an indication of the reliability or trust in the final estimate.
- 2) If reliability is high, the fused estimate is more accurate than any of the individual ones.
- 3) If reliability is limited, the system generates an explanation in human terms.
- 4) The fused estimate is more robust.

These characteristics allow the user to determine the suitability of the estimate within the given business application context. Knowledge-based rules are used for constructing this fusion of technologies at a supervisory level, and the

parameters are few enough to be determined by some inspection and experimentation.

B. Related Work

Combining forecasts and estimates has been extensively studied in disciplines such as econometrics and market research [117]–[121]. A common fusion technique is to use a weighted average where the weights are related to some measure of reliability in the individual forecasts. This is essentially the approach we have taken in this work. On the other hand, we believe the method we have used to obtain a combined reliability is unique.

C. Approach to Fusion

Our approach is broken down into four steps.

- 1) Measure the reliability or “goodness” of each individual estimate. This is done individually by each estimator, using a fuzzy definition of “typicality.”
- 2) Combine the three dollar values produced by the individual estimators into a fused dollar value.
- 3) Use a rule base that determines the overall quality of the fused estimate—this is referred to as the fusion of the reliability values.
- 4) Generate explanatory messages to the user when the overall quality is limited. These messages enable the user to understand why the best possible estimate cannot be produced. Further information by the user may then enable the system to improve its results.

We will refer to the overall system as ARPV. The final output of the system will be referred to as the *ARPV-Estimate*. The accompanying measure of reliability is called the *Quality* of this estimate.

D. Fusion Rules

As seen in Fig. 10, LOCVAL is used as an input to AIGEN, and AICOMP and AIGEN both use property descriptions to produce the answers. These two are then combined to produce the final *ARPV-Estimate* of property value and an overall measure of reliability or *Quality*. Finally, ARPV provides a description of circumstances for the data, property, and environment that may have led to a limited value of *Quality*.

The precise manner of their combination depends on the relative strengths of the two individual reliability values, and the amount of disagreement between the two estimator values. Each reliability value is a number in [0,1]. The amount of disagreement is measured by *contention* between the two values v_{AIG} (produced by AIGEN), and v_{AIC} (produced by AICOMP). Contention c is computed as

$$c = (v_{AIG} - v_{AIC}) / ((v_{AIG} + v_{AIC}) / 2)$$

and is defined only when both estimators have produced a valid answer.

ARPV-Estimate, the fused property value, takes one of the following values.

- 1) The mean of the two estimates, weighted by their respective normalized reliability values—if both AIGEN and AICOMP estimates are known.
- 2) The value of the available estimator—if one of AIGEN and AICOMP estimates is known but the other is not, making their combination impossible.
- 3) The value of LOCVAL—if both AIGEN and AICOMP are unknown, making their combination impossible.

Quality, the measure of the reliability of the fused ARPV estimate, takes one of the following three values.

- 1) EXCELLENT—All essential data are available, reliability for the independent estimators is high, and the independent estimators agree.
- 2) INDICATIVE—All essential data are available, reliability of one or more estimators is not high due to somewhat unusual property and/or local market characteristics, and/or the estimators disagree.
- 3) UNRELIABLE—Some essential data may be missing, and/or reliability in one or more estimators is low due to very unusual property and/or local market characteristics, and/or the estimators disagree markedly.

A more detailed description of the fusion rules can be found in [116].

E. Fusion Results

The reliability value generated by the fusion rules was subdivided into three groupings (good, fair, and poor). The reliability measure should then produce the largest good set with the lowest error.

From a test sample of 7293 subjects, we could classify our reliability in 63% as good. The good set had a medium absolute error of 5.4%, an error that was satisfactory for

the intended application. Of the remaining subjects, 24% were classified as fair, and 13% as poor. The fair set had a medium error of 7.7%, and the poor set had a median error of 11.8%.

F. Summarizing the SC Applications to Property Valuation

We have considered the important and difficult problem of residential property valuation and shown that techniques based on SC can successfully solve it. A multitude of approaches using FL, case-based reasoning, and NN's (in hybrid combination) were shown to be useful in this regard. Moreover, the reliability computation and the fusion process increase the robustness and human usefulness of the system. It has been shown to achieve good accuracy and be scalable for thousands of automated transactions. This makes it a transparent, interpretable, fast, and inexpensive choice for bulk estimates of residential property value for a variety of financial applications.

IX. FINAL REMARKS

SC is having an impact on many industrial and commercial operations, from predictive modeling to diagnostics and control. It provides us with alternative approaches to traditional knowledge-driven reasoning systems or pure data-driven systems (the fundamental problems of these classical approaches are discussed in Section I-D) and it overcomes their shortcomings by synthesizing a number of complementary reasoning and searching methods over a large spectrum of problem domains. We have reviewed SC's main components (FL, probabilistic reasoning, NN's, and EC) and surveyed some of their successful combinations that have led to the development of hybrid SC systems.

These systems leverage the tolerance for imprecision, uncertainty, and incompleteness, which is intrinsic to the problems to be solved, and generate tractable, low-cost, robust solutions to such problems. The synergy derived from these hybrid systems stems from the relative ease with which we can translate problem domain knowledge into initial model structures whose parameters are further tuned by local or global search methods. This is a form of complementary or "tight" hybridization. Apart from this type of hybridization, we also discussed the fusion of estimators—this type of model fusion or "loose" hybridization does not combine features of the methodologies themselves, but only their results. The primary motivation here is to increase reliability rather than to make model construction easier.

We have illustrated this synergy by describing three applications in diagnostics, control, and prediction, as summarized in Table 3. These studies stem from real-world, high-impact business problems, such as gas turbine service and diagnosis, freight train control, and residential property valuation, respectively, and they use a number of different means from the tool box of SC. In particular, all approaches used a hybridization of SC techniques to overcome hurdles posed by the problem, such as high noise and low infor-

Table 3
SC Applications in Industry and Business Discussed in This Paper

Application Area	Problem Type	Type of Hybrid	Model Structure Type	Parameter Search Method
Industrial equipment anomaly detection	Diagnostics	Complementary	Fuzzy clusters in high-dimensional space	Exponential Weighted Moving Average Filter
Automatic train handling	Nonlinear Control	Complementary	Fuzzy PI Controller rule-base	Genetic algorithms
Residential Property Valuation	Predictive Modeling	Complementary	Fuzzy if-then rules as network	Variation on back propagation / gradient descent
Residential Property valuation	Predictive Modeling	Fusion	Rule-base for fusing result & reliability	Manually determined

mation content of sensor data, high system complexity, and high-dimensional feature space.

The payoff of this conjunctive use of techniques is a more accurate and robust solution than a solution derived from the use of any single technique alone. This synergy comes at comparatively little expense because typically the methods do not try to solve the same problem in parallel but they do it in a mutually complementary fashion. Another way to say this is that the model needs a structure and parameters, and a search method to discover them, and no single technique should be expected to be the best for all problems. For example, in our control application, a hierarchical fuzzy controller was used to embody the qualitative knowledge used by locomotive engineers to manually perform the task, but manual tuning of this controller would have been an inferior solution. Fast tuning of this controller was obtained by a global search approach, a GA that yielded a robust controller whose parameters did not have to be specialized for each particular initial condition.

Another advantage to the hybridization of techniques is that it is easier to think of alternative solutions to the same problem, as was evidenced by the property estimation application. If there are several possibilities for the structure and the search methods, many more pairings of technologies are possible, and problem solving becomes easier. For instance, AIGEN used a fuzzy model with local gradient search, but one could also use a pure NN trained by an evolutionary algorithm. Of course, computation time, cost, business needs, and data requirements will then influence the choice of the combination of technologies. A step in further improving system performance is the exploitation of parallel systems. These systems may be designed to rely to the maximum amount on nonoverlapping data and use different techniques to arrive at their conclusions. In information fusion, the outputs of these heterogeneous models will be compared, contrasted, and aggregated, as seen in our last application.

The future appears to hold a lot of promise for the novel use and combinations of SC applications. The circle of SC-related technologies will probably widen beyond its current constituents. The push for low-cost solutions combined with the need for intelligent tools will result in the deployment of hybrid systems that efficiently integrate reasoning and search techniques.

On the application front we will likely see a drive toward prognostic and autonomous capabilities. With an increase of service-related operations, it will be increasingly attractive to be able to forecast anomalous trends and conditions, and correct them before their effects are fully developed. In addition, remotely monitored systems will bear the need to operate autonomously, thus requiring intelligent agents to regulate their operations.

In the future, we expect that the combination of SC with advances in the areas of computer vision, voice recognition, natural language processing, etc., will further improve and expand our problem-solving capability to a large spectrum of industrial and commercial problems.

REFERENCES

- [1] P. P. Bonissone, V. Badami, K. H. Chiang, P. S. Khedkar, K. Marcelle, and M. J. Schutten, "Industrial applications of fuzzy logic at General Electric," *Proc. IEEE*, vol. 83, pp. 450–465, Mar. 1995.
- [2] Y.-T. Chen and P. P. Bonissone, "Industrial applications of neural networks at General Electric," General Electric CRD, Schenectady, NY, Tech. Inform. Series 98CRD79, Oct. 1998.
- [3] P. P. Bonissone and W. Cheatham, "Financial applications of fuzzy case-based reasoning to residential property valuation," in *Proc. 6th Int. Conf. Fuzzy Systems (FUZZ-IEEE'97)*, Barcelona, Spain, 1997, pp. 37–44.
- [4] L. A. Zadeh, "Fuzzy logic and soft computing: Issues, contentions and perspectives," in *Proc. IIZUKA'94: 3rd Int. Conf. Fuzzy Logic, Neural Nets and Soft Computing*, Iizuka, Japan, 1994, pp. 1–2.
- [5] —, "Some reflection on soft computing, granular computing and their roles in the conception, design and utilization of information/intelligent systems," *Soft Comput. Fusion of Foundations, Methodologies Applicat.*, vol. 2, no. 1, pp. 23–25, 1998.
- [6] D. Dubois and H. Prade, "Soft computing, fuzzy logic, and Artificial Intelligence," *Soft Comput. Fusion of Foundations, Methodologies Applicat.*, vol. 2, no. 1, pp. 7–11, 1998.
- [7] B. Bouchon-Meunier, R. Yager, and L. A. Zadeh, *Fuzzy Logic and Soft Computing*. Singapore: World Scientific, 1995.
- [8] P. P. Bonissone, "Soft computing: The convergence of emerging reasoning technologies," *Soft Comput. Fusion of Foundations, Methodologies Applicat.*, vol. 1, no. 1, pp. 6–18, 1997.
- [9] N. Rescher, *Many-Valued Logic*. New York: McGraw-Hill, 1969.
- [10] M. Black, "Vagueness: An exercise in logical analysis," *Phil. Sci.*, vol. 4, pp. 427–455, 1937.
- [11] L. A. Zadeh, "Fuzzy sets," *Inform. Control*, vol. 8, pp. 338–353, 1965.
- [12] J. Lukasiewicz, *Elementy Logiki Matematycznej (Elements of Mathematical Logic)*. Warsaw, Poland: Panstwowe Wydawnictwo Naukowe, 1929.
- [13] L. A. Zadeh, "Foreword," in *Handbook of Fuzzy Computation*, E. H. Ruspini, P. P. Bonissone, and W. Pedycz, Eds. Bristol, U.K.: Inst. Phys., 1998.

- [14] E. H. Ruspini, P. P. Bonissone, and W. Pedycz, *Handbook of Fuzzy Computation*. Bristol, U.K.: Inst. Phys., 1998.
- [15] Y.-M. Pok and J.-X. Xu, "Why is fuzzy control robust," in *Proc. 3rd IEEE Int. Conf. Fuzzy Systems (FUZZ-IEEE'94)*, Orlando, FL, 1994, pp. 1018–1022.
- [16] T. Bayes, "An essay toward solving a problem in the doctrine of chances," *Philosophical Trans. Roy. Soc. London*, vol. 53, pp. 370–418, 1763; see also *Facsimiles of Two Papers by Bayes*. New York: E. Deming, 1963.
- [17] E. Shortliffe and B. Buchanan, "A model of inexact reasoning in medicine," *Math. Biosci.*, vol. 23, nos. 3–4, pp. 351–379, 1975.
- [18] R. Duda, P. Hart, and N. Nilsson, "Subjective Bayesian methods for rule-based inference systems," in *Proc. AFIPS*, vol. 45. New York, NY: AFIPS, 1976, pp. 1075–1082.
- [19] J. Pearl, "Reverend Bayes on inference engines: A distributed hierarchical approach," in *Proc. 2nd Nat. Conf. Artificial Intelligence*. Menlo Park, CA: AAAI, 1982, pp. 133–136.
- [20] ———, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann, 1988.
- [21] J. Kim and J. Pearl, "A computational model for causal and diagnostic reasoning in inference engines," in *Proc. 8th Int. Joint Conf. Artificial Intelligence*, Karlsruhe, Germany, 1983, pp. 190–193.
- [22] A. P. Dempster, "Upper and lower probabilities induced by a multivalued mapping," *Ann. Math. Statist.*, vol. 38, pp. 325–339, 1967.
- [23] G. Shafer, *A Mathematical Theory of Evidence*. Princeton, NJ: Princeton Univ. Press, 1976.
- [24] L. A. Zadeh, "Probability measures of fuzzy events," *J. Math. Anal. Appl.*, vol. 10, pp. 421–427, 1968.
- [25] P. Smets, "The degree of belief in a fuzzy set," *Inform. Sci.*, vol. 25, pp. 1–19, 1981.
- [26] W. S. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bull. Math. Biophys.*, vol. 5, pp. 115–133, 1943.
- [27] F. Rosenblatt, "The perceptron, a perceiving and recognizing automaton," Cornell Aeronautical Lab. Rep., Buffalo, NY, Project PARA, no. 85-640-1, 1957.
- [28] ———, "Two theorems of statistical separability in the perceptron," in *Proc. Mechanization of Thought Processes*, London, U.K., 1959, pp. 421–456.
- [29] F. Rosenblatt, *Principle of Neurodynamics: Perceptron and the Theory of Brain Mechanisms*. Washington, DC: Spartan, 1962.
- [30] M. Minsky and S. Papert, *Perceptrons*. Boston, MA: MIT Press, 1969.
- [31] P. Werbos, "Beyond regression: New Tools for predictions and analysis in the behavioral science," Ph.D. dissertation, Harvard Univ., Cambridge, MA, 1974.
- [32] D. Parker, "Learning logic," Center for Computational Res. Econ. Manage. Sci., Massachusetts Inst. Technol., Cambridge, MA, Tech. Rep. TR-47, 1985.
- [33] Y. LeCun, "Une procedure d'apprentissage pour reseau a seuil symetrique," *Cognitiva*, vol. 85. Paris, France: CESTA, 1985, pp. 599–604.
- [34] K. Hornick, M. Stinchcombe, and H. White, "Multilayer feed-forward networks are universal approximators," *Neural Networks*, vol. 2, pp. 359–366, 1989.
- [35] J. Moody and C. Darken, "Fast learning in networks of locally tuned processing units," *Neural Computation*, vol. 1, no. 2, pp. 281–294, 1989.
- [36] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Bio. Cybern.*, vol. 43, pp. 59–69, 1982.
- [37] J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," in *Proc. Acad. Sci.*, vol. 79, pp. 2554–2558, 1982.
- [38] A. Carpenter and S. Grossberg, "A Massively parallel architecture for a self-organizing neural pattern recognition machine," *Comput., Vision, Graphics, Image Processing*, vol. 37, pp. 54–115, 1983.
- [39] R. Jang, C.-T. J. Sun, and C. Darken, "Functional equivalence between radial basis function networks and fuzzy inference systems," *IEEE Trans. Neural Networks*, vol. 4, pp. 156–159, Jan. 1993.
- [40] E. Fiesler and R. Beale, *Handbook of Neural Computation*. New York: Oxford, 1997.
- [41] I. Rechenberg, "Cybernetic solution path of an experimental problem," *Royal Aircraft Establishment*, Library Translation no. 1122, 1965.
- [42] H.-P. Schwefel, "Kybernetische evolution als strategie der experimentellen forschung in der stromungstechnik," Diploma thesis, Tech. Univ. Berlin, Germany, 1965.
- [43] L. J. Fogel, "Autonomous automata," *Ind. Res.*, vol. 4, pp. 14–19, 1962.
- [44] L. J. Fogel, A. J. Owens, and M. J. Walsh, *Artificial Intelligence through Simulated Evolution*. New York: Wiley, 1966.
- [45] A. S. Fraser, "Simulation of genetic systems by automatic digital computers, I: Introduction," *Australian J. Bio. Sci.*, vol. 10, pp. 484–491, 1957.
- [46] H. Bremermann, "The evolution of intelligence. The nervous system as a model of its environment," Dep. Math., Univ. Washington, Seattle, Tech. Rep. 1, Contract 477(17), 1958.
- [47] J. Reed, R. Tooms, and N. Baricelli, "Simulation of biological evolution and machine learning," *J. Theo. Bio.*, vol. 17, pp. 319–342, 1967.
- [48] J. H. Holland, "Outline of a logical theory of adaptive systems," *J. ACM*, vol. 9, pp. 297–314, 1962.
- [49] ———, "Nonlinear environments permitting efficient adaptation," *Computer and Information Science IIs*. New York: Academic, 1967.
- [50] ———, *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press, 1975.
- [51] J. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press, 1992.
- [52] R. M. Friedberg, "A learning machine: Part I," *IBM J. Res. Develop.*, vol. 3, pp. 282–287, 1958.
- [53] N. A. Barricelli, "Esempi numerici di processi di evoluzione," *Methodos*, pp. 45–68, 1954.
- [54] D. B. Fogel, *Evolutionary Computation*. New York: IEEE Press, 1995.
- [55] T. Bäck, D. B. Fogel, and Z. Michalewicz, Eds., *Handbook of Evolutionary Computation*. New York: Oxford, 1997.
- [56] D. B. Fogel, *The Fossil Record*. New York: IEEE Press, 1998.
- [57] I. Rechenberg, *Evolutionstrategien: Optimierung Technischer Systeme nach Prinzipien der Biologischen Evolution*. Stuttgart, Germany: Fromman-Holzboog Verlag, 1973.
- [58] H.-P. Schwefel, *Numerical Optimization of Computer Models*. Chichester, U.K.: Wiley, 1981.
- [59] P. J. Angeline, G. M. Saunders, and J. B. Pollack, "An evolutionary algorithm that constructs recurrent neural networks," *IEEE Trans. Neural Networks*, vol. 5, pp. 54–65, Jan. 1994.
- [60] Z. Michalewicz, *Genetic Algorithms + Data Structures = Evolution Programs*. New York: Springer-Verlag, 1994.
- [61] K. Forbus, "Qualitative reasoning about physical processes," in *Proc. 7th Int. Joint Conf. Artificial Intelligence*, Karlsruhe, Germany, 1981.
- [62] B. Kuipers, "Commonsense reasoning about causality: Deriving behavior from structure," in *Qualitative Reasoning about Physical Systems*, D. Bobrow, Ed. Cambridge, MA: MIT Press, 1985, pp. 169–203.
- [63] E. H. Mamdani and S. Assilian, "An experiment in linguistic synthesis with a fuzzy logic controller," *Int. J. Man Machine Studies*, vol. 7, no. 1, pp. 1–13, 1975.
- [64] T. Takagi and M. Sugeno, "Fuzzy identification of systems and its applications to modeling and control," *IEEE Trans. Syst., Man, Cybern.*, vol. 15, pp. 116–132, Jan.–Feb. 1985.
- [65] R. Babuska, R. Jager, and H. B. Verbruggen, "Interpolation issues in Sugeno-Takagi reasoning," in *Proc. 3rd IEEE Int. Conf. Fuzzy Systems (FUZZ-IEEE'94)*, Orlando, FL, 1994, pp. 859–863.
- [66] H. Bersini, G. Bontempi, and C. Decaestecker, "Comparing RBF and fuzzy inference systems on theoretical and practical basis," in *Proc. Int. Conf. Artificial Neural Networks. ICANN '95*, vol. 1, Paris, France, 1995, pp. 169–174.
- [67] J. S. R. Jang, "ANFIS: Adaptive-network-based-fuzzy-inference-system," *IEEE Trans. Syst., Man, Cybern.*, vol. 23, pp. 665–685, May–June 1993.
- [68] O. Cordon, H. Herrera, and M. Lozano, "A classified review on the combination fuzzy logic-genetic algorithms bibliography," Dep. Comput. Sci. AI, Univ. Granada, Granada, Spain, Tech. Report 95129, 1995. [Online]. Available WWW: <http://decsai.ugr.es/herrera/flga.html>.
- [69] C. L. Karr, "Design of an adaptive fuzzy logic controller using genetic algorithms," in *Proc. Int. Conf. Genetic Algorithms (ICGA'91)*, San Diego, CA, 1991, pp. 450–456.
- [70] M. A. Lee, and H. Tagaki, "Dynamic control of genetic algorithm using fuzzy logic techniques," in *Proc. 5th Int. Conf. Ge-*

- netic Algorithms. San Mateo, CA: Morgan Kaufmann, 1993, pp. 76–83.
- [71] H. Surmann, A. Kanstein, and K. Goser, "Self-organizing and genetic algorithms for an automatic design of fuzzy control and decision systems," in *Proc. EUFIT'93*, Aachen, Germany, 1993, pp. 1097–1104.
- [72] J. Kinzel, F. Klawoon, and R. Kruse, "Modifications of genetic algorithms for designing and optimizing fuzzy controllers," in *Proc. 1st IEEE Conf. Evolutionary Computing (ICEC'94)*, Orlando, FL, 1994, pp. 28–33.
- [73] D. Burkhardt and P. P. Bonissone, "Automated fuzzy knowledge base generation and tuning," in *Proc. 1st IEEE Int. Conf. Fuzzy Systems (FUZZ-IEEE'92)*, San Diego, CA, 1992, pp. 179–188.
- [74] P. P. Bonissone, P. S. Khedkar, and Y.-T. Chen, "Genetic algorithms for automated tuning of fuzzy controllers, a transportation application," in *Proc. 5th IEEE Int. Conf. Fuzzy Systems (FUZZ-IEEE'96)*, New Orleans, LA, 1996, pp. 674–680.
- [75] L. Zheng, "A practical guide to tune proportional and integral (PI) like fuzzy controllers," in *Proc. 1st IEEE Int. Conf. Fuzzy Systems, (FUZZ-IEEE'92)*, San Diego, CA, 1992, pp. 633–640.
- [76] V. W. Port, "Overview of Evolutionary Computation as a mechanism for solving neural system design problems," *D2.1 in Handbook of Neural Computation*, E. Fiesler and R. Beale, Eds. New York: Oxford, 1997.
- [77] E. Vonk, L. C. Jain, and R. P. Johnson, *Automatic Generation of Neural Network Architecture Using Evolutionary Computation*. Singapore: World Scientific, 1997.
- [78] X. Yao "Evolving artificial neural networks," this issue, pp. 1423–1447.
- [79] M. Jurick, "Back error propagation: A critique," *Proc. IEEE COMPCON 88*, San Francisco, CA, 1988, pp. 387–392.
- [80] D. J. Montana and L. Davis, "Training feedforward neural networks using genetic algorithms," in *Proc. 11th Int. Joint Conf. Artificial Intelligence (IJCAI-89)*, vol. 1. San Francisco, CA: Morgan Kaufmann, pp. 762–767, 1989.
- [81] H. Kitano, "Empirical studies on the speed of convergence of neural network training using genetic algorithms," in *Proc. 8th National Conf. Artificial Intelligence (AAAI-90)*, vol. 2, 1990, pp. 789–795.
- [82] M. McInerney and A. P. Dhawan, "Use of genetic algorithms with backpropagation in training of feedforward neural networks," in *Proc. 1993 IEEE Int. Conf. Neural Networks (ICNN '93)*, vol. 1, San Francisco, CA, 1993, pp. 203–208.
- [83] V. Maniezzo, "Genetic evolution of the topology and weight distribution of neural networks," *IEEE Trans. Neural Networks*, vol. 5, pp. 39–53, Jan. 1994.
- [84] S. Harp, T. Samad, and A. Guha, "Toward the genetic synthesis of neural networks," in *Proc. 3rd Int. Conf. Genetic Algorithms*, 1989, pp. 360–369.
- [85] E. Alba, J. F. Aldana, and J. M. Troya, "Genetic algorithms as heuristics for optimizing ANN design," in *Proc. Int. Conf. Artificial Neural Nets and Genetic Algorithms (ANNGA'93)*, Innsbruck, Austria, 1993, pp. 683–689.
- [86] H. Kitano, "Designing neural networks using genetic algorithms with graph generation system," *Complex Syst.*, vol. 4, no. 4, pp. 461–476, 1990.
- [87] A. A. Siddiqi and S. M. Lucas, "A comparison of matrix rewriting versus direct encoding for evolving neural networks," in *Proc. IEEE World Congr. Computational Intelligence (WCCI'98)*, Anchorage, Alaska, 1998, pp. 392–397.
- [88] F. Gruau, "Genetic synthesis of Boolean neural networks with a cell rewriting developmental process," in *Proc. 2nd Int. Conf. Genetic Algorithms (COGAN'92)*, Baltimore, MD, 1992, pp. 55–74.
- [89] J. Koza, and J. P. Rice, "Genetic generation of both the weights and architecture for a neural network," in *Proc. IEEE Int. Joint Conf. Neural Networks*, 1991, pp. 397–404.
- [90] D. B. Fogel, L. J. Fogel, and V. W. Porto, "Evolutionary methods for training neural networks," in *Proc. IEEE Conf. Neural Networks for Oceanography Engineering*, Washington, DC, 1991, pp. 317–327.
- [91] J. R. McDonnell and D. Waagen, "Evolving neural network connectivity," in *Proc. IEEE ICNN'93*, San Francisco, CA, pp. 863–868.
- [92] D. E. Goldberg, "Genetic algorithms and Walsh functions: Part 2, deception and analysis," *Complex Syst.*, vol. 3, pp. 153–171, 1989.
- [93] D. B. Fogel, "Evolving artificial intelligence," Ph.D. dissertation, Univ. California, San Diego, 1992.
- [94] C.-H. Lee and J.-H. Kim, "Evolutionary ordered neural network with a linked-list encoding scheme," in *Proc. IEEE ICEC'96*, pp. 665–669.
- [95] R. A. Jacobs, "Increased rates of convergence through learning rate adaptation," *Neural Networks*, vol. 1, pp. 295–307, 1988.
- [96] P. Arabshahi, J. J. Choi, R. J. Marks, and T. P. Caudell, "Fuzzy control of backpropagation," in *Proc. 1st IEEE Int. Conf. Fuzzy Systems (FUZZ-IEEE'92)*, San Diego, CA, 1992, pp. 967–972.
- [97] P. D. Wasserman, *Neural Computing: Theory and Practice*. New York: Van Nostrand Reinhold, 1989.
- [98] F. Herrera and M. Lozano, "Adaptive genetic algorithms based on fuzzy techniques," in *Proc. IPMU'96*, Granada, Spain, 1996, pp. 775–780.
- [99] R. Subbu, A. Anderson, and P. P. Bonissone, "Fuzzy logic controlled genetic algorithms versus tuned genetic algorithms: An agile manufacturing application," in *Proc. IEEE Int. Symp. Intelligent Control, NIST*, Gaithersburg, MD, 1998.
- [100] D. Doel, "The role for expert systems in commercial gas turbine engine monitoring," in *Proc. Gas Turbine and Aeroengine Congr. Exposition*, Brussels, Belgium, 1990.
- [101] S. Karamouzis and S. Feyock, "A performance assessment of a case-based diagnostic system for aircraft malfunctions," in *Proc. 6th Int. Conf. Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, Edinburgh, U.K., 1993, pp. 71–78.
- [102] R. Reibling and S. Bublin, "Diagnostic reasoning technology for the on-board maintenance system," in *Proc. IEEE 1993 National Aerospace and Electronics Conf. NAECON 1993*, vol. 2, pp. 930–936.
- [103] M. Fernandez-Montesinos, P. Janssens, and R. Vingerhoeds, "Enhancing aircraft engine condition monitoring. Safety, reliability and applications of emerging intelligent control technologies," *Proc. IFAC Workshop Emerging Intelligent Control Technologies*, Hong Kong, 1994, pp. 161–166.
- [104] R. Records and J. Choi, "Spurious symptom reduction in fault monitoring using a neural network and knowledge as hybrid system," in *Proc. 12th National Conf. Artificial Intelligence*, vol. 2, Seattle, WA, 1994.
- [105] J. B. Gomm, "Process-fault diagnosis using a self-adaptive neural network with on-line learning capabilities," in *Proc. IFAC On-Line Fault Detection and Supervision in the Chemical Process Industry*, 1995, pp. 69–74.
- [106] V. C. Patel, V. Kadirkamanathan, and H. A. Thompson, "A novel self-learning fault detection system for gas turbine engines," in *Proc. UKACC Int. Conf. Control '96*, pp. 867–872.
- [107] E. H. Ruspini, "A new approach to clustering," *Inform. Control*, vol. 15, no. 1, pp. 22–32, 1969.
- [108] J. J. Grefenstette, "GENESIS: A system for Using Search Procedures," in *Proc. 1984 Conf. Intelligent Systems and Machines*, pp. 161–165.
- [109] K. A. De Jong, *An Analysis of the Behavior of a Class of Genetic Adaptive Systems*, Ph.D. dissertation, Univ. Michigan, Ann Arbor, 1975.
- [110] G. Rudolph, "Convergence of evolutionary algorithms in general search spaces," in *Proc. IEEE Int. Conf. Evolutionary Computation*, Nagoya, Japan, 1996, pp. 50–54.
- [111] Appraisal Institute, *Appraising Residential Properties, Part VI*, Chicago, IL, 1994.
- [112] J. Kolodner, *Case-Based Reasoning*. San Francisco, CA: Morgan Kaufmann, 1993.
- [113] Breiman, Friedman, Olshen, and Stone, *Classification and Regression Trees*. Monterey, CA: Wadsworth and Brooks, 1985.
- [114] A. J. Gonzalez, "A case-based reasoning approach to real estate property appraisal," *Expert Syst. Applicat.*, vol. 4, no. 2, pp. 229–246, 1992.
- [115] P. P. Bonissone and S. Ayub, "Similarity measures for case-based reasoning systems," in *Proc. 4th Int. Conf. Information Processing and Management of Uncertainty (IPMU-92) in Knowledge-Based Systems*, Palma, Spain, 1992, pp. 483–487.
- [116] P. P. Bonissone, W. Cheetham, D. Golibersuch, and P. S. Khedkar, "Automated residential property valuation: An accurate and reliable based on soft computing," in *Soft Computing in Financial Engineering*, R. Ribeiro, H. Zimmermann, R. R. Yager, and J. Kacprzyk, Eds. Heidelberg, Germany: Physica-Verlag, 1998.
- [117] D. Bunn, "Combining forecasts," *Europ. J. Oper. Res.*, vol. 33, no. 3, pp. 223–229, 1988.

- [118] R. T. Clemen, "Combining forecasts: A review and annotated bibliography," *Int. J. Forecasting*, vol. 5, no. 4, pp. 559–584, 1989.
- [119] C. W. J. Granger and R. Ramanathan, "Improved methods of combining forecasts," *J. Forecasting*, vol. 3, pp. 197–204, 1984.
- [120] E. Kacapyr, "Consensus forecasts," *Economic Forecasting: The State of the Art*. New York: M. E. Sharpe, 1996.
- [121] S. K. McNees, "Consensus forecasts: Tyranny of the majority," *New England Econ. Rev.*, pp. 14–21, Nov.-Dec. 1987.



Piero P. Bonissone (Member, IEEE) received the B.S. degree in electrical engineering and mechanical engineering from the University of Mexico City, Mexico, in 1975 and the M.S. degree in electrical engineering and computer science (EECS), the M.S. degree in mechanical engineering, and the Ph.D. degree in EECS from the University of California, Berkeley, in 1976, 1978, and 1979, respectively.

He has been a Computer Scientist at the General Electric Corporate Research and Development Center (GE CRD), Niskayuna, NY, since 1979. There he has carried out research and projects in artificial intelligence, expert systems, simulation, fuzzy sets, soft computing, and data mining. In the early 1980's, he led the development team for Diesel Electric Locomotive Troubleshooting Aid (DELTA), one of the first fielded expert systems, which helped maintenance technicians in troubleshooting diesel electric locomotives. He was the Principal Investigator in the research program on Reasoning with Incomplete and Uncertain Information, as part of the Knowledge Based Systems Technology Base in DARPA's Strategic Computing Program. He was also responsible for the development of the situation assessment modules in DARPA's Pilot's Associate and Submarine Operational Automation System (SOAS) projects. He led the development of the Reasoning with Uncertainty Module (RUM) and PRIMO, two General Electric (GE) proprietary expert system tools for plausible reasoning. He has led many projects in fuzzy logic control and soft computing, ranging from the control of turbo-shaft engines to the use of fuzzy logic in dishwashers, locomotives, and power supplies. He has also developed case-based and fuzzy-neural systems to accurately estimate the value of single family residential properties when used as mortgage collaterals. In 1990, he developed combined approximate reasoning systems (CARS), a case-based reasoning tool integrated with PRIMO and applied to military transportation planning problems. Recently he has led a soft computing (SC) group in the development of SC applications to diagnostics and prognostics of processes and products. He has been an Adjunct Professor of the ECSE Department at Rensselaer Polytechnic Institute (RPI), Troy, NY, since 1982. Since 1993, he has been the Editor-in-Chief of the *International Journal of Approximate Reasoning*. He co-edited the books *Expert Systems in Structural Safety Assessment* (Springer, 1989), *Uncertainty in Artificial Intelligence 6* (North-Holland, 1991), and *Uncertainty in Artificial Intelligence 7* (Morgan Kaufmann, 1991). He is also an Editor-in-Chief of the *Handbook of Fuzzy Computation (Institute of Physics)*. He has published more than 90 articles in the area of expert systems, approximate reasoning, fuzzy sets, pattern recognition, decision analysis, and SC. He received seven patents from the U.S. Patent Office for his work on reasoning with uncertainty and fuzzy control.

Dr. Bonissone received the Dushman Award from GE CRD in 1989 for his work on reasoning with uncertainty. In 1993, he received the Coolidge Fellowship Award from GE CRD for overall technical accomplishments. In 1996 he became a Fellow of the American Association for Artificial Intelligence (AAAI). In 1999, he received the Dushman Award from GE CRD for his work on medical equipment diagnostics.



Yu-To Chen (Member, IEEE) received the B.S. degree in mechanical engineering from National Central University, Taiwan, in 1985 and the M.S. and the Ph.D. degrees in industrial engineering with a minor in computer science from the Pennsylvania State University, University Park, in 1990 and 1993, respectively.

From 1993 to 1994, he was a Visiting Assistant Professor at the University of Iowa. He is currently a Computer Scientist in the Information Technology Laboratory, General Electric Corporate Research and Development Center (GE CRD), Niskayuna, NY. His research interests include data mining, soft computing, and statistics as applied to industrial and financial systems.

Dr. Chen is a member of the IIE. He is also a recipient of the Dushman Award, the most prestigious technical team award at GE CR&D, in 1999.



Kai Goebel received the Diplom-Ingenieur degree from the Technische Universität München, Munich, Germany, in 1990 and the M.S. and Ph.D. degrees in mechanical engineering from the University of California (UC), Berkeley in 1993 and 1996, respectively.

He joined the General Electric Corporate Research and Development Center (GE CRD), Niskayuna, NY, in 1997 as a Computer Scientist after working as a Visiting Postdoctoral Fellow at UC Berkeley from 1996 to 1997. He has carried out applied research in the areas of artificial intelligence, soft computing, and information fusion. He has worked on using soft computing techniques for real-time monitoring and diagnosis of aircraft engines, power plants, and manufacturing processes, and for both data fusion and decision fusion in automated vehicle control systems, aircraft engines, and manufacturing systems. He has been an Adjunct Professor of the Computer Science Department at Rensselaer Polytechnic Institute (RPI), Troy, NY, since 1998 where he teaches classes in soft computing.



Pratap S. Khedkar received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, Delhi, India, in 1982 and the Ph.D. degree in computer science from the University of California, Berkeley, in 1993.

Since 1993, he has been a Computer Scientist at the General Electric Corporate Research and Development Center (GE CRD), Niskayuna, NY. While there, he has carried out applied research in the areas of artificial intelligence (AI), neural networks, fuzzy systems, expert systems, soft computing, and data mining. He has worked on using AI technology for the control and management of locomotives, paper mills, polymer manufacturing, and electric grid relays. His recent work focuses on the application of data mining and soft computing technology to problems in the financial services domain, ranging from risk analysis, property valuation, target marketing, pricing, and reinsurance. He has also served as an Adjunct Professor in the ECSE Department at Rensselaer Polytechnic Institute (RPI), Troy, NY. He was an Associate Editor of the *Handbook of Fuzzy Computation (Institute of Physics)*.