In this paper I present a design process model called GENCAD (GENetic Case ADaptation) that combines ideas from CBR and GA's. The CBR paradigm provides a method for the overall process of case selection and adaptation. The GA paradigm provides a method for adapting design cases by combining and mutating their features until a set of new design requirements and constraints are satisfied. The requirements of a new design problem are specified by the user at the outset of the process. Domain-imposed design constraints can serve as additional knowledge for evaluating the quality of the solutions proposed through the adaptation of previously known designs by the process.

The GENCAD process model has been implemented computationally. The type of design problem used to illustrate the model is *feng shui* residential floor plan layout. I use a set of Frank Lloyd Wright prairie house layouts as the case base. Problem requirements are specified by the user as the number of rooms of different types that are required in a new design. The constraints used to determine whether new designs proposed by the process model are acceptable are taken from *feng shui*, the Chinese art of placement. This implementation not only illustrates how the process model for design through the evolutionary adaptation of cases works, but it also shows how different knowledge sources with distinct origins can be used within the same design framework.

In section 2 of the paper I present GENCAD, the process model for the evolutionary adaptation of design cases. In section 3 I discuss the sample application domain, *feng shui* residential layout design. In section 4 I present the results of two experiments I performed with the implementation of GENCAD to the *feng shui* domain. Section 5 closes the paper by summarizing and discussing the material presented in the other sections.

2 GENCAD: A Process Model for the Evolutionary Adaptation of Design Cases

The GENCAD process model emphasizes solving problems using information derived from precedents; this is the main idea behind CBR. The solutions to past design problems contained in the precedents that are retrieved from a case memory serve as starting points for proposing solutions to new design problems. Multiple random combinations and modifications (together referred to as adaptations) of the retrieved cases are then generated and evolved incrementally, until a satisfactory solution to the new design problem is found; this is the main idea behind GA's. Figure 1 shows the GENCAD process model for the evolutionary adaptation of design cases.

Given the specification of the requirements of a new design problem, the first task in the process model is to determine which precedents contain information that might be useful in solving the new problem. This is done by consulting a case base and comparing the description of the new problem with the descriptions of the precedents stored in the case base. Those cases for which some similarity is found with the new problem are retrieved from the case base. They are taken to be first approximations towards a solution to the new problem, and are put together into a population of potential designs. The cases in this population then need to be adapted to become satisfactory solutions to the new problem.

In GENCAD this task of case adaptation is performed by an evolutionary method. Two types of adaptation are available to the evolutionary method: combination and modification. These types of adaptation are performed on the designs in the population through the genetic operators of crossover and mutation, respectively. Both crossover and mutation involve random decisions; thus, the precise appearance of the new designs resulting from adapting the old designs cannot be predicted before performing the adaptations. Crossover produces two offspring designs, each of which combines features from each of two parent designs. Mutation produces one offspring design which is an altered version of one parent design. Both crossover and mutation insert new designs into the population.

The two types of adaptation can result in generating offspring designs that are better (i.e., closer to being a solution to the new design problem) than some or all of the old parent designs, or worse. In order to determine this, the potential solutions have to be evaluated and their relative worth compared. In the process of evaluating the solutions, one or more might be found that are good enough to satisfy the requirements of the new design problem. If this happens, the process model can stop at this point; if not, the evolutionary adaptation process continues, using the best of the designs as the initial population for the next cycle of adaptations. The best designs are selected from the augmented population containing both old (relative to the current evolutionary cycle) and newly generated potential designs.

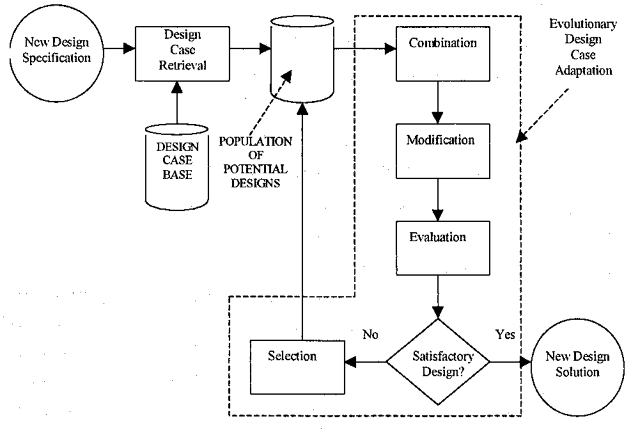

Figure 1. GENCAD: a process model for case-based design with evolutionary case adaptation

Evaluation and selection are not directly responsible for manipulating old designs to produce new ones, unlike combination and modification. However, they are considered part of the evolutionary case adaptation process because the results they give define the paths examined by GENCAD during the search for a solution to the new problem. Evaluation and selection also define which designs will be available for adaptation in future cycles; thus, they guide and influence case adaptation.

The GENCAD process model is generic; that is, it is applicable to a wide variety of design domains, and makes no commitments as to how design cases, phenotypes, and genotypes will be represented. However, in order to prove its feasibility I wrote a LISP implementation of it, applied to the domain of *feng shui* residential layout design. The following section describes this domain briefly.

3 *Feng Shui* Residential Layout Design

The domain chosen as proof-of-concept for testing the GENCAD process model is the design of residential layouts that satisfy the principles of *feng shui*. This section describes some of the characteristics of this domain and some knowledge representation and processing issues raised by it.

3.1 Characteristics of *feng shui*

Feng shui, also known as Chinese geomancy, is an ancient Chinese technique that, among other things, determines the quality of proposed or existing layouts of residences according to several rules of thumb. Some of these heuristics seem to have a basis in common sense, or in a psychological or sociological appreciation of the human beings that inhabit (or intend to inhabit) the residence. Other heuristics seem to be of a more superstitious nature.

There are several different *feng shui* sects that may contradict each other or place different priorities on different aspects of residential layouts. Some *feng shui* heuristics focus on the date of birth, profession, gender, income, and family relationships of the people that the residence is or will be for. They might interpret the quality of a design according to a combination of these social characteristics with some geographical features such as absolute cardinal directions (north, south, etc.). Other *feng shui* heuristics might emphasize more the presumed speed at which *ch'i* (positive energy assumed to emanate from living creatures and animate objects) is allowed to circulate or flow by a particular design.

Despite this variety, of prime importance to performing a *feng shui* analysis according to the principles of any of the existing sects is information on the relative positions of objects. Some other attributes of objects that are usually also taken into account are their orientations, shapes, and relative sizes. In the work reported in this paper I have used the knowledge of *feng shui* presented in (Rossbach 1987), which corresponds to the Tibetan black-hat sect of *feng shui*.

Feng shui analyses different aspects of a residential layout to determine its auspiciousness or lack thereof. Some classes of inauspicious layouts can be "cured" by the proper placement of an acceptable curing object. Thus, some potentially bad layouts can actually be acceptable if the proper cure is present. Given a layout, it is not just a matter of determining whether the layout is "good" or whether something about it is "bad," because even if some aspect of it makes it bad, one has to determine whether a valid cure for its bad features is present or not before rejecting the layout outright.

The *feng shui* knowledge contained in (Rossbach 1987) applies to three different levels of description of a residence:

- the landscape level (the location of a residence with respect to other objects in its environment such as mountains, rivers, roads, etc.),
- the house level (the relative placement of the rooms and functional spaces within a residence, such as bedrooms and bathrooms, as well as the connections between them, such as doors and windows), and
- the room level (the location of furniture, decorations, and other objects within each room or functional space in a residence).

Feng shui analysis assumes knowledge of spatial relationships among the objects at each of the different levels. Absolute locations and exact measures of distances and other geometric quantities are not as important. Because of this, a qualitative spatial representation of location seems to be sufficient.

3.2 Spatial representation for *feng shui*

The representational framework I used to describe the locations of objects within each of the three levels is a low-resolution, qualitative one. I locate objects (which I will refer to as elements) at each level on a 3x3 spatial grid, with each sector within the grid assigned a unique number between 1 and 9 to identify it. The grid is shown in Figure 2, with north assumed to be at the top of the grid.

1	2	3
4	5	6
7	8	9

Figure 2. Qualitative spatial grid for elements

Elements can occupy more than one grid sector, and grid sectors can contain more than one element, making the representation flexible. The resolution of this representation is not high, but considering the qualitative nature of a typical *feng shui* analysis and the number of elements that typically need to be represented at each of the three levels, it is adequate in most cases. It allows the determination of the relative positions of objects, which is the most important capability needed for an adequate *feng shui* analysis.

Figure 3 shows a typical residence viewed at the three levels of description that a *feng shui* analysis would look at, with the 3x3 spatial grid superimposed on each level. The figure is immediately followed by an equivalent symbolic representation of the residence shown in it. This demonstrates how the relevant information can be represented in the computer.

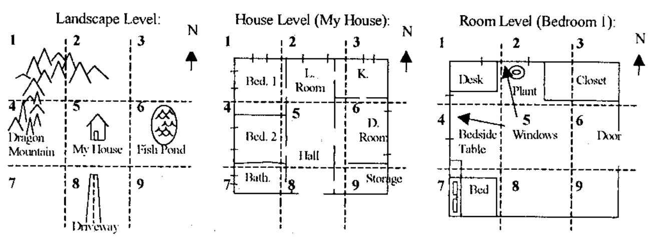

Figure 3. Example of a residence viewed at the three levels

```
(((level landscape)
  (elements (((type hill) (name dragon-mountain)
              (location (1 2 4)) (steepness high) ...)
             ((type house) (name my-house)
              (location (5))) ...)))
 ((level house)
  (elements (((type bedroom) (name bedroom-1)
              (location (1))   (shape square))
             ((type bedroom)  (name bedroom-2)
             (location (4)))
            ((type hallway) (name hall-1)
             (location (5 8))) ...))
  (connectors (((type internal-door)
                (name b2-hall) (location (5))
                 (side-a bedroom-2) (side-b hall-1)
                (direction ew))
               ((type window) (name b1-window-1)
                 (location (1)) (side-a bedroom-1)
                (side-b outside) (direction ew)) ...)))
 ((level room)
  (name bed-1)
  (elements (((type bed) (name b-1)
              (location (7)))
       ((type desk) (name d-1)
        (location (1)))
       ((type window) (name w-1)
       (location (1 2))) ...))))
```

3.3 *Feng shui* Analysis

Constraints can be used for evaluation of individuals in the context of GENCAD's evolutionary method of case adaptation. In the *feng shui* domain these constraints implement what can be termed *feng shui* analysis, i.e., analyzing potential solutions generated by the evolutionary method according to how well they satisfy the principles of *feng shui*.

An example of a constraint at the house level is given by the following quote from (Rossbach 1987):

> Traditionally, the Chinese avoid three or more doors or windows in a row...this...funnels ch'i [positive energy] too quickly...[CURE:]...to stop ch'i from flowing too quickly, hang a wind chime or crystal ball... (page 89)

This constraint is implemented by first finding the description of all the connectors (doors and windows) at the house level, particularly focusing on their locations and directions. If at least three connectors are aligned such that their locations are in consecutive (or the same) grid sectors *and* they all have the same direction (e.g., north-south), then the constraint has been violated. However, before determining this it must be determined whether or not there are any wind chimes or crystal balls in the house that are positioned in line with the violating doors/windows. The algorithm that represents this constraint procedurally is the following (given a phenotype P):

```
Get the list C of all connectors in P;
Get the list Q of all potential cures in P for this constraint;
For each connector c in C, or until a bad omen has been
found:
  Get the location l of c;
  Get the direction d of c;
  Set the list of connectors LU lined up with c
  to NIL;
  Get the list Reduced of all elements in C except c;
  For each connector r in Reduced:
    If the direction of r is d And
      the location of r lines up with l along
      direction d,
      Then add r to LU;
    End-If;
  End-For;
  If there are two or more connectors in LU And
    no potential cure in Q lines up with r,
    Then signal a bad omen situation;
  End-If;
End-For;
```

4 Experimental Results

Applying the GENCAD process model to the domain of design of residences such that they satisfy the principles of *feng shui* resulted in a system named GENCAD-FS. The programming of GENCAD-FS introduced the need to choose representational frameworks for design cases, phenotypes, genotypes, and *feng shui* constraints, taking into account the specific characteristics of this domain. Details of the representational and implementational intricacies of GENCAD-FS can be found in (Gómez de Silva Garza and Maher, 1999a) and (Gómez de Silva Garza and Maher, 1999b). This section reports on two experiments performed with GENCAD-FS.

4.1 Experiment I: Solution Quality

In one experiment I performed with GENCAD-FS the purpose was to analyze the quality of the solutions found by comparing their features to those stylistic characteristics normally present in the Frank Lloyd Wright cases that were used as a basis for constructing the solutions (which were obtained from (Hildebrand, 1991)).

The following house-level problem specifications were given in order to perform this experiment:

((bedroom 3) (bathroom 2) (fireplace 1) (music-room 1))

That is, I asked GENCAD-FS to propose the floor plan of a residence with (at least) three bedrooms, two bathrooms, a fireplace, and a music room, such that the relative positions of these rooms and architectural features satisfy the principles of *feng shui*. The solution to this problem that was found by GENCAD-FS is shown in Figure 4.

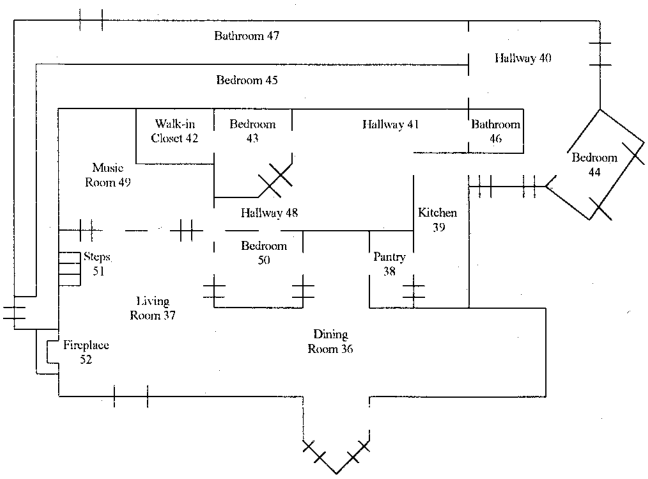

Figure 4. Solution found by GENCAD-FS

Some characteristics of the process that GENCAD-FS underwent to produce the solution in Figure 4, as well as of the solution itself, are the following:

- All 12 Frank Lloyd Wright cases were retrieved from memory in order to participate in adaptation, as they all matched one or more of the example problem requirements.
- GENCAD-FS converged after 61 evolutionary cycles to the solution shown.
- A total of 28 crossover operations and 3 mutations took place in order to generate this solution (not counting the crossover and mutation operations whose results eventually caused a dead end during the search).
- The solution combines features inherited from three of the twelve cases that were originally retrieved: Avery Coonley House, Robie House (main floor), and Willits House (lower floor). Features from the other nine retrieved cases were probably still present in the population at the 61^{st} generation, but not in the final solution (and were also never part of any of its ancestors).

Analyzing the floor plan shown in Figure 4, and taking into account the observations on this solution made above, the following conclusions about the inheritance of individual features and the creativity of the resulting design can be reached:

- The fact that the solution is a floor plan that contains many hallways is a feature that is inherited from Avery Coonley House, which has 4 hallways. Neither of the other two ancestor cases has more than two hallways.
- The protuberance from the dining room is a feature that is inherited from Robie House, which has two such similar protuberances (one from the dining room and one from the living room), or from Willits House which has one (from the dining room).
- The room that is skewed compared to the north-south grid (Bedroom 44) is a feature that is inherited from Willits House, in which the rooms are all similarly skewed.
- The bedrooms in the solution are all in very different parts of the house, something that does not normally occur in a residence. In a family house, at least, the bedrooms are usually all near each other. However, if the residence were intended for several unrelated individuals sharing a house, the solution would be a creative and interesting design that would allow the occupants to have greater privacy than the design of a typical family house.

In other words, the solution combines (in a creative way) features from several known previous solutions. The solution also has some of the features considered to be general aspects of Frank Lloyd Wright's style, such as having large dining and living rooms that merge into each other, with no partitions between them, and having a fireplace. The fireplace was a problem requirement, so any solution had to have one to be a valid solution; its location in the living room was inherited from one of the three ancestor cases (all three had this feature).

4.9 Experiment II: Problem Solving Efficiency

Another experiment served to test the efficiency of GENCAD's proposed method of design problem solving. This experiment was performed by giving GENCAD-FS the following landscape-level specification:

((house 1) (river 1) (trees 2))

This problem specification can be interpreted as "I want to build a house on a property in which there is a river, and I'm thinking of planting two clumps of trees around the house." The problem is now to use GENCAD-FS to generate a configuration containing these four elements, specifying their relative positions within the landscape, such that the configuration is auspicious according to the principles of *feng shui.*

In order to perform the efficiency experiment, GENCAD-FS was run 20 times using the 12 cases from its Frank Lloyd Wright prairie house case base (they are all retrieved, given the above problem requirements), and 20 times using 12 randomly-generated "cases," on the same problem. Using the 12 random "cases" to initialize the non-case-based evolutionary search is equivalent to seeding the initial population of a GA randomly. The 12 random "cases" only include landscape-level information, as this is the only type of information needed for the experiment. The cases were generated by randomly choosing how many landscape-level elements to include in them (between 1 and 10). For each of these elements, its type was randomly selected, the value of its location attribute was randomly assigned, and depending on the type of the element, the values of other possible relevant attributes were randomly assigned. For the assignment of values, a global variable embodying knowledge of which valid values exist for each possible attribute for each type of element was used.

In this experiment, any difference in efficiency can be attributed to the use of CBR as the guiding framework. In this way I can evaluate the decision to combine the two paradigms of CBR and GA's, as opposed to using a traditional GA by itself (without cases) to search for solutions, which could have been an alternative way of solving the same problems.

GENCAD-FS was given a limit of 500 GA cycles in which to find an acceptable solution. If convergence did not occur by cycle 500, the search was ended without a solution being given. Some of the cases in the randomly generated case base, as well as the Frank Lloyd Wright cases, do contain two clumps of trees, and/or a house, and/or a river in the landscape. In addition, there are configurations of these four types of element that are valid according to *feng shui* practice. Therefore, achieving a solution through the cyclical combination and/or modification of either of the two sets of cases is theoretically possible for GENCAD-FS.

In the experiment, 5 of the 20 trials using the random starting points converged. Similarly, 5 of the 20 trials using the Frank Lloyd Wright cases converged. Thus, whether cases or random starting points are used to initiate the search doesn't seem to make a difference as far as the frequency of convergence (of course, further experiments would need to be done to support this claim more strongly). However, a clear difference can be seen when we analyze the number of GA cycles required before convergence occurred (in those trials in which it did occur), as seen in Table 1.

Trial #	Random	Trial #	FLW cases
1	114	25	54
9	333	31	34
11	357	36	32
14	274	37	406
17	160	39	90
Avg.:	241.6	**Avg.:**	123.2

Table 1. GA cycles required before convergence

As can be seen from the results, when cases are used to guide (i.e., to provide starting points for) the evolutionary search, convergence occurs on average twice as fast as when the search is initiated from random starting points. This demonstrates the efficiency of combining the ideas of CBR with those from GA's. As can also be seen, convergence does not always occur, or does not occur within a reasonable number of iterations. Whether it will converge or not, or how rapidly it will converge, can vary greatly due to the random nature of the genetic operators of crossover and mutation. However, as the experiment also shows, the GENCAD process model can be applied again and again to the same problem, using the same initial set of retrieved cases, and it is possible that it will converge in future attempts even if it didn't at first. This is something that might not be possible with "pure" CBR systems which are deterministic and therefore cannot continue trying to find a solution to the same problem by trying again once they get stuck.

5 Summary and Discussion

I have presented GENCAD, a process model for creative design that combines aspects of case-based reasoning and evolutionary algorithms. From the point of view of case-based reasoning, the evolutionary algorithm in GENCAD performs the task of case adaptation. In order to do so it requires domain knowledge with which to recognize good solutions, but does not require domain knowledge with which to generate potential adaptations. This simplifies the implementation of the case adaptation task compared to "traditional" case-based reasoning systems. From the point of view of evolutionary algorithms, the cases retrieved by GENCAD given a new problem specification serve as an initial population of potential solutions for the algorithm to operate on. This gives the evolutionary algorithm an initial bias/help that "traditional" evolutionary algorithms, which are normally initialized with a population of randomly-generated potential solutions, don't have. GENCAD benefits from the advantages of both case-based reasoning and evolutionary algorithms and avoids their respective disadvantages by combining the two methods in such a way that both methods end up helping each other.

The application of GENCAD to the design of the floor plan layout of residences such that the designs produced conform to the principles of *feng shui* was used to illustrate the process model. One of the things this illustration showed is the opportunity that GENCAD provides for combining knowledge originating from diverse design communities, something I have termed *inter-cultural design*. Inter-cultural design can be used in process models to produce design solutions that contain interesting combinations of features. A process model for design with this characteristic can be used to simulate, enhance, or explore human creativity.

References

Gómez de Silva, A., and **Maher, M.L.** (1999a). An Evolutionary Approach to Case Adaptation. In Althoff, K.D., Bergmann, R., and Branting, L.K. (eds.) *Case-Based Reasoning Research and Applications: Proceedings of the Third International Conference on Case-Based Reasoning ICCBR-99*, 162-172. Berlin: Springer-Verlag.

Gómez de Silva, A., and **Maher, M.L.** (1999b). Evolving Design Layout Cases to Satisfy Feng Shui Constraints. *Proceedings of the Fourth Conference on Computer Aided Architectural Design in Asia CAADRIA-99*, 115-124. Shanghai: Shanghai Scientific and Technological Literature Publishing House.

Hildebrand, G. (1991). *The Wright Space: Pattern & Meaning in Frank Lloyd Wright's Houses*. Seattle: University of Washington Press.

Leake, D.B. (1996). *Case-Based Reasoning: Experiences, Lessons, & Future Directions*. Menlo Park, California, and Cambridge, Massachusetts: AAAI Press and MIT Press.

Maher, M.L., and **Pu, P.** (eds). (1997). *Issues and Applications of Case-Based Reasoning in Design*. Mahwah, New Jersey: Lawrence Erlbaum Associates.

Mitchell, M. (1998). *An Introduction to Genetic Algorithms (Complex Adaptive Systems Series)*. Cambridge, Massachusetts: MIT Press.

Rossbach, S. (1987). *Interior Design with Feng Shui.* London: Rider Books.

Andrés Gómez de Silva Garza *received his undergraduate degree in Computer Engineering from the National Autonomous University of México (UNAM) in 1991. That same year he participated for six months in an IBM Systems Training Internship. Afterwards, in 1994, he got his Master's degree from the Georgia Institute of Technology in Atlanta. He then obtained his Ph.D. from the University of Sydney in Australia in 2000, for which his research was on applying multiple Artificial Intelligence techniques to design tasks. He is currently a full-time lecturer in the Computer Science Department of the Mexican Autonomous Technological Institute (ITAM).*

Mary Lou Maher *obtained her undergraduate degree in Civil Engineering from Columbia University in 1979. She then obtained both her Master's degree (1981) and her Ph.D. (1984) from Carnegie-Mellon University, both involving the use of computers for aiding in design and engineering tasks. She has previously been on the academic staff at Carnegie-Mellon and is now a full Professor and Co-Director of the Key Centre of Design Computing and Cognition of the University of Sydney, Australia. She has authored, co-authored, or edited more than one hundred books, book chapters, journal papers, and articles on her research interests.*