Financial Applications of Fuzzy Case-Based Reasoning to Residential Property Valuation

Piero P. Bonissone
GE R&D Center
One Research Circle
Niskayuna, NY 12309, USA
Bonissone@crd.ge.com

William Cheetham
GE R&D Center
One Research Circle
Niskayuna, NY 12309, USA
Cheetham@crd.ge.com

Abstract

We present PROFIT, an advanced prototype system developed to estimate residential property values for real estate transactions. The system enhances Case-Based Reasoning (CBR) techniques with fuzzy predicates expressing preferences in determining similarities between subject and comparable properties. These similarities guide the selection and aggregation process, leading to the final property value estimate. Fuzzy techniques are also used to generate a confidence value qualifying such estimate. PROFIT has been successfully tested on thousands of real estate transactions.

1 Introduction

Residential property valuation [1] is the process of determining a dollar estimate of the property value for given market conditions. A residential property is 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 (appraisal) must be periodically updated to reflect those market changes. Any appraisal must also be supported by current evidence of market conditions, e.g. recent real estate transactions.

1.1 Problem Description and Motivation

Many financial institutions grant mortgages and purchase mortgage packages on the secondary market as investments. These packages can contain up to 1000 mortgages. Appraisals are needed to grant most new mortgages and to evaluate the current value of mortgage packages that may be purchased. The current manual process for appraising properties usually requires an on-site visit by a human appraiser, lasts three to four days, and costs about $500 per subject property. This process is too slow and expensive for packages of 1000 mortgages, whose value is currently estimated, to a lesser degree of accuracy, by sampling techniques. This manual process is also subject to human variability, as appraisers of varying experience will provide appraisals of varying accuracy. Some appraisers may even try to validate a negotiated price rather than estimating the true value of a property. This practice is generally easier but less accurate than estimating the true market value of the property.

The most common method used by human appraisers is referred to as 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 reasonable averaging method). This process assumes[1] that the item's market value can be derived by the prices demanded by similar items in the same market.

The possibility of automating this process was first shown by Gonzalez [4]. However, his Case-Based Reasoner (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.

To address this problem we have developed the PROperty Financial Information Technology (PROFIT) system, which enhances case-based reasoning techniques [7] with fuzzy predicates and fuzzy-logic based similarity measures [2] to estimate the value of residential property. Our approach consists of:

1) Retrieving recent sales from a case-base using a small number of features to select potential comparables.

2) Comparing the subject property with the retrieved cases and deriving a partial ordering (similarity measure) from the aggregation of fuzzy preference values.

3) Adjusting the sales price of the retrieved cases to reflect their differences from the subject using a rule set.

4) Aggregating the adjusted sales prices of the retrieved cases, selecting the best comparables, deriving a single value estimate for the subject, and qualifying the estimate with a confidence value.

[1] It is also assumed that sales comparable are instances of arm-length transactions of willing buyers and sellers in a reasonably efficient market.

0-7803-3796-4/97/$10.00©1997IEEE

1.2 Paper Structure

We will briefly present some related works and competitive approaches followed by a description of our proposed fuzzy CBR process. We will then illustrate the process with an example, providing an estimate value for the subject property and computing a confidence value for the estimate. Finally, we will present a statistical analysis of the overall process, followed by a summary of our findings, conclusions, and future work.

2 Related Works

In a team effort with other researchers at GE, we developed three other (non CBR based) methods for estimating the value of a property. The first alternative approach was based only on the location of comparable properties. All properties within 0.1 mile of the subject were combined, giving the closer properties higher weight, to produce a typical unit price for living area [$/sq.ft] for any property in that location. The unit price was multiplied by the subject's living area to produce an estimate. The second method used a statistical approach involving classification and regression trees [3] to produce a formula, based on ten attributes of the subject, and generate an estimate of the subjects value. The third method trained a fuzzy-neural net (ANFIS) [5] using a subset of cases from the case-base, and produced a run-time system to provide an estimate of the subjects value. Figure 1 describes the data requirements, method used, and median of the absolute relative error for these competitive approaches, and compares them with the fuzzy CBR approach and the human appraiser. We can observe that the methods' accuracy increases monotonically with their data requirements.

Data Needed	Method Used	Error(Median)
2 attributes	$/sq. ft. in area	10%
10 attributes	Statistical formula	8%
10 attributes	Fuzzy-Neural Net	7%
[11-30] attributes	Fuzzy CBR	5%
Site Inspection	Human Appraiser	3%

Figure 1: Comparison of multiple approaches

3 Fuzzy CBR Process

PROFIT consists of a case-base of every property that has been sold in California during the last five years and a process for selecting relevant cases, adapting them, and aggregating those adapted cases into a single estimate of the property value. The case-base[2] was stored in a Sybase database, containing several hundred thousands sale records of California real estate transactions. The records were purchased from the California Market Data Corporation (CMDC) and TRW. These data vendors provided a description of each property which, in some cases, contained up to 166 property attributes.

The property valuation process is shown in Figure 2. Upon entering the subject property attributes, PROFIT retrieves potentially similar comparables from the case-base. This initial selection uses six attributes: address, date of sale, living area, lot area, number of bathrooms, and bedrooms. 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 making preferences, expressed as fuzzy membership distributions and relations. Each property's sales price is adjusted to better reflect the subject's value. These adjustments are performed by a rule set that uses additional property attributes, such as construction quality, conditions, pools, fireplaces, etc. The best 4 to 8 comparables are then selected. Finally, the adjusted sales price and similarity of the selected properties are combined to produce an estimate of the value of the subject, a confidence in that estimate, and a justification for the estimate.

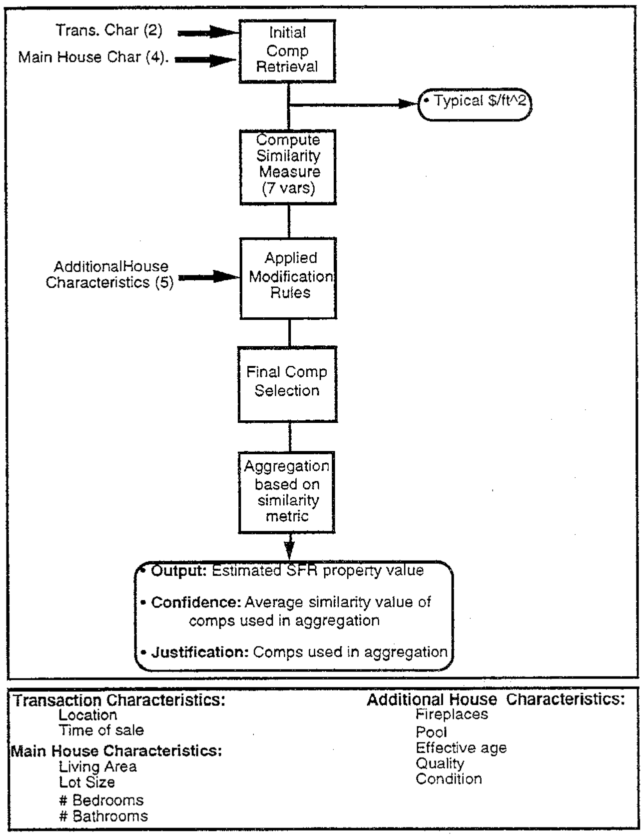

Figure 2: PROFIT Process

4 Example

We will now provide an example of the property valuation process, consisting of the following steps:

[2] Customized data integrity and filtering techniques were needed to transform the original database into a suitable case-base.

- Retrieval of comparables from a database
- Evaluation of each comparable's attribute on a preference scale from 0 to 1
- Computation of each comparable's similarity with the subject, (weighted sum of preference values)
- Adjustment of comparables sale prices to reflect differences from the subject property
- Removal of poor quality comparables
- Weighted aggregation of best comparables

4.1 Case Retrieval

The initial retrieval extracts a set of potential comparables using standard SQL queries for efficiency purpose. The selection is performed by comparing specific attributes of the subject with the corresponding attribute of each comparable. All the comparables in the retrieved set have values within the allowable deviations. If the size of the retrieved set is too small, e.g., less than 10, the allowable deviations could be relaxed to increase its at the expense of retrieval quality.

This initial retrieval stage uses the following attributes and their corresponding maximum allowable deviations (written after each attribute):

- Date of sale (within 12 months)
- Distance (within 1 mile)
- living area (+ / - 25%)
- lot size (+ 100% / - 50%)
- Number of bedrooms (+/- 3)
- Number of bathrooms (+/- 3)

These ranges correspond to the support of the fuzzy sets shown in Figure 3 and the fuzzy relations of Figures 4-5.

Rationale for attribute selection. After performing a completeness analysis on CMDC and TRW data bases, we decided to use the above six key attributes because their fields contained a value in over 95% of the records in the data bases. The first two attributes (number of months since the date of sale, and distance from subject) are market and region dependent. Their range of allowed values could be manually modified or automatically indexed to reflect slow or fast markets, as well as urban, suburban, and rural regions. The remaining four variables (living area, lot area, number of bedrooms, number of bathrooms) reflect some of the subject's main characteristics. The range of values used in the initial retrieval stage have been obtained from several knowledge engineering sessions with two GE appraisers.

4.2 Preference Criteria Definition and Evaluation.

Figure 3 describes our preference criteria for the first four features. The trapezoidal membership distributions representing these criteria have a natural preference interpretation. For each feature, the *support* of the distribution represents the range of ***tolerable*** values and corresponds to the interval-value used in the initial retrieval query. The *core* represents the most ***desirable*** range of values and establishes our top preference. By definition, a feature value falling inside the core will receive a preference value of 1. As the feature value moves away from the most desirable range, its associated preference value will decrease from 1 to 0. At the end of this evaluation, each comparable will have a ***preference vector***, with each element taking values in the (0,1] interval. These values represent the partial degree of membership of each feature value in the fuzzy sets and fuzzy relations representing our preference criteria.

For example, by using the preference distributions shown in Figure 3 we can see that the preference value for the attribute *date-of-sale* of a comparable that was sold within 3 months of today's date is 1. If the date was 6 months ago, its preference value would be 2/3. Any comparable with a date of sale of more than 12 months would be given a preference value of zero.

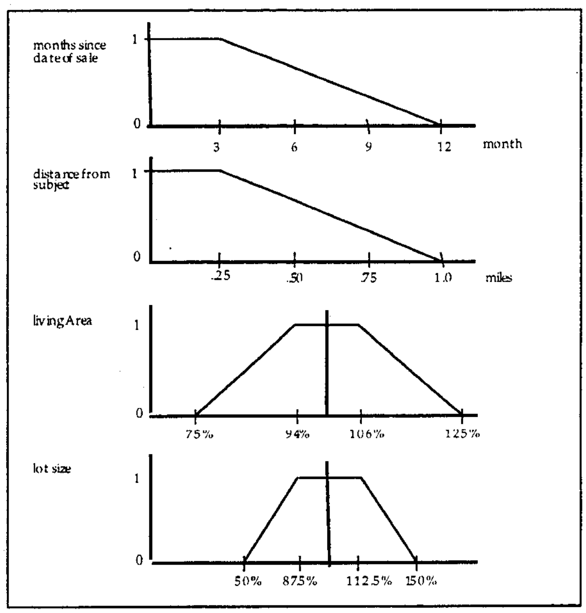

Figure 3: Attribute Preference Functions

The remaining two features, *Number of Bedrooms* and *Number of Bathrooms,* are evaluated in a similar fashion. Their preference functions are represented by two reflexive asymmetric fuzzy relations, illustrated in Figure 4 and 5, respectively.

Comparable's # Bedrooms		1	2	3	4	5	6 +
	1	1.00	0.50	0.05	0.00	0.00	0.00
Subject's	2	0.20	1.00	0.50	0.05	0.00	0.00
# Bedrooms	3	0.05	0.30	1.00	0.60	0.05	0.00
	4	0.00	0.05	0.50	1.00	0.60	0.20
	5	0.00	0.00	0.05	0.60	1.00	0.80
	6 +	0.00	0.00	0.00	0.20	0.80	1.00

Figure 4: Preference Function for Number of Bedrooms

For instance, using Figure 4, we can observe that for a subject with 5 bedrooms the preferred comparable would also have 5 bedroom (preference =1), while a 6 bedroom comparable would meet that preference criterion to a degree of 0.8. Similarly, using Figure 5, we can observe that, for a subject with 2 bathrooms, the preferred comparable would also have 2 bathrooms (preference value =1), while a 2.5 bathroom comparable would meet that preference criterion to a degree of 0.7.

Subject	Comparable								
	1	1.5	2	2.5	3	3.5	4	4.5	5+
1	1.00	0.75	0.20	0.05	0.01	0.00	0.00	0.00	0.00
1.5	0.60	1.00	0.60	0.25	0.10	0.05	0.00	0.00	0.00
2	0.10	0.70	1.00	0.70	0.25	0.05	0.00	0.00	0.00
2.5	0.05	0.20	0.75	1.00	0.75	0.20	0.05	0.00	0.00
3	0.01	0.10	0.40	0.80	1.00	0.80	0.40	0.10	0.05
3.5	0.00	0.05	0.15	0.45	0.85	1.00	0.85	0.45	0.30
4	0.00	0.00	0.05	0.20	0.50	0.90	1.00	0.90	0.70
4.5	0.00	0.00	0.00	0.10	0.30	0.70	0.95	1.00	0.95
5+	0.00	0.00	0.00	0.05	0.15	0.35	0.75	0.95	1.00

Figure 5: Preference Function for Number of Bathrooms

4.3 Similarity Measure Computation: Preference Weighting and Aggregation

The next step consists in computing a similarity measure between each potential comparable and the subject. The similarity measure is a function of the preference vector computed above and of the decision maker's priorities. These priorities are reflected by the weights used before the aggregation. Figure 6 (sixth column) illustrates a set of values obtained by interviewing expert appraisers, using Saaty's pairwise comparison method [9], and validated in our tests.

Figure 6 shows an example of the similarity measure computation between the subject and a comparable. Attributes (living area, lot area, bedrooms, bathrooms) and derived values (date, distance) used for the calculation are listed in the column labeled *Attribute*. The values of those attributes for the subject and comparable are listed in the columns labeled *Subject* and *Comparable*. The evaluation representing the attributes' degree of matching, obtained by using Figure 3, 4, and 5 are listed in the *Preference* column. The weights reflecting an attribute's relative importance in the specific market area are listed in the *Weight* column. The sum of the score and weights for an attribute is given in the *Weighted Preference* column. Finally, the total score of the weighted preferences represents the similarity measure.

4.4 Comparables Adjustments.

All of the properties found by the initial retrieval will undergo a series of adjustment in their sales price to better reflect the subject property value. Any difference between the subject and comparable property that would cause the comparable to be more (or less valuable) than the subject produces an adjustment. If the comparable is superior to the subject that adjustment will decrease the comparable price, and viceversa for comparable that are inferior to the subject. After all the adjustments are applied to the comparable sales price the resulting value is called the comparable *Adjusted price*.

The adjustment rules, illustrated in Table 1, were obtained from numerous knowledge engineering sessions with expert appraiser and from the analysis of hundreds of existing appraisals. These rules will be triggered by differences between subject's and comparable's attributes.

Living Area	*(subject - comp) * (22 + (Sales_Price_of_comp * .00003))*
Lot Area	*(subject - comp) * 1*
Bathrooms	*see figure 7*
Fireplaces	*(subject - comp) * 2000*
Effective Year Built	
*w * (Age_comp-Age_subject) * (Sale_Price_comp/1000)* *if (Age_subject + Age_comp) / 2 < 4 then w = 4 else* *if (Age_subject + Age_comp) / 2 < 6 then w = 3 else* *if (Age_subject + Age_comp) / 2 < 8 then w = 2 else* *if (Age_subject + Age_comp) / 2 < 15 then w = 1 else* *w = .5* *max of 10% of salePrice*	
Quality	*(.02 * sale price) for each level of difference:*
(Luxury > Excellent > Good> Average > Fair > Poor)	
Pool	*$10000 for a pool*

Table 1: Adjustment Rule Set

Figure 7 shows the adjustments (in thousands of dollars) to be made to the comparable's price, as a function of the different number of bathrooms between the subject and the comparable property. The last column and row indicate the required adjustments for each additional bathroom when either the subject or the comparable have more than five bathrooms.

Figure 8 shows the computation of the comparable's adjusted price. The adjustments shown in the figure are:

- The Living Area is adjusted by (22 + (175000 * .00003)) = $27.25 per square foot which is 200 * $27.25 = $5450.
- The Lot Area is adjusted by $1/sq ft for a total of -$5000 since the comparable has a larger lot size.
- The Bathrooms are adjusted using Figure 7 (adjustment figures in 1,000 of dollars).
- There are the same number of Bedrooms so there is no adjustment.
- The subject has one more Fireplace than the comparable so the adjustment is $2000.
- The Effective Year Built formula produces an adjustment of $2800.
- The subjects Quality is one step better than the comparable so the adjustment is 2% of the sale price of the comparable.
- The subject has a Pool and the comparable does not so the adjustment is $10000.

Attribute	Subject	Comparable	Comparison	Preference	Weight	Weighted Preference
Months since date of sale	X	6 months	6 months	0.67	0.222	0.1489
Distance	X	0.2 miles	0.2 miles	1.00	0.222	0.2222
Living Area	2000	1800	90%	0.79	0.333	0.2633
Lot Size	20000	35000	175%	0.33	0.111	0.0367
# Bedroom	3	3	0%	1.00	0.056	0.0556
# Bathrooms	2.5	2	2.5 -> 2	0.75	0.056	0.0417
Similarity Measure (Sum of Weighted Preference/Sum of Weights) =						**0.768333**

Figure 6: Similarity Measure Computation

Subject	Comp								
	1	1.5	2	2.5	3	3.5	4	4.5	5+
1	**0.00**	*-1.50*	-3.00	*-5.00*	*-8.00*	N/A	N/A	N/A	N/A
1.5	1.00	**0.00**	-1.00	*-3.50*	*-6.00*	*-9.00*	N/A	N/A	N/A
2	*4.00*	*1.50*	**0.00**	-2.25	*-4.00*	*-6.50*	N/A	N/A	N/A
2.5	*7.00*	*4.50*	2.00	**0.00**	-2.00	*-4.50*	*-7.00*	N/A	N/A
3	*9.00*	*6.50*	3.00	2.00	**0.00**	-2.50	-5.00	*-7.50*	'@*-5
3.5	N/A	*8.50*	*6.50*	*4.50*	*2.50*	**0.00**	*-3.00*	*-5.50*	'@*-5
4	N/A	N/A	*8.50*	*7.00*	*5.50*	*3.00*	**0.00**	*-3.00*	'@*-5
4.5	N/A	N/A	N/A	*10.00*	*8.00*	*6.00*	*3.00*	**0.00**	'@*-5
5+	N/A	N/A	N/A	'@*5	'@*5	'@*5	'@*5	'@*5	**0.00**

Figure 7: Adjustment Function for Number of Bathrooms

Attribute	Subject	Comparable	Adjustment
SalePrice	?	175000	175000
LivingArea	2000	1800	5450
LotArea	20000	25000	-5000
SFRTotalBaths	2.5	2	2000
SFRBedrooms	3	3	
SFRFireplaces	1	0	2000
EffYearBuilt	93	89	2800
Quality	Good	Average	3500
Condition	Average	Average	
Pool	Yes	No	10000
Adjusted Price =			**195750**

Figure 8: Example of Adjustments

4.5 Comparables Filtering

A good comparable selection is key to the property valuation process. We have found that the range of four to eight comparables is optimal for this process. By using less than four comparables, we are not correctly reflecting the current market. By using more than eight, we risk to include some comparables which are not similar enough to the subject property. If it is not possible to find four comparables similar to the subject property then no value estimate may be calculated for the subject. Typically, the initial retrieval yields an average of 22 comparables (up to a maximum of 100). Therefore, finding the best 4 - 8 comparables will usually require removing the less desirable comparables. We would like the selected comparables to have the following properties:

- No single adjustment should be larger (in absolute value) than 10% of sales price
- Net adjustment should not exceed 15% of sales price
- Gross adjustment should not exceed 25% of sales price
- The unit price for living area of the comparables should not vary more than 15% from each other and should bracket that of the subject
- Comparables should be as close as possible to the subject
- The value estimated for the subject should be bracketed by the sales price of the comparables

After adjusting all the comparables found in the initial retrieval, any comparable violating any of the first three constraints is removed from the process.

The best comparables are selected from the remaining ones by the method show in Figure 9. First, we determine the comparables with the highest similarity score, lowest net adjustments, and lowest gross adjustments. Figure 10 shows how these three values are combined into a single ranking of the comparables. Each comparable has its similarity score ranked with the other comparables such that the comparable with the best similarity score receives the lowest rank. Net adjustment and gross adjustment are similarly ranked. The sum of the three ranks is computed for each comparable, determining its *Total Rank.* The comparables with the lowest total rank are considered the best.

This method induces a cardinal ordering on the set of comparables, but does not tell us how many comparables to use. The sales prices of the comparables should bound the estimated sales price of the subject. Therefore, we would like to select comparables with both negative and positive net adjustments. Typically, a comparable with a negative net adjustment is likely to have an unadjusted price greater than the final estimate, and viceversa for a comparable with a positve net adjustment. To achieve this bracketing effect, we create a temporary set of candidates by repeatedly adding the comparable with the best similarity score to the set until: 1) there are at least four comparables in the set and 2) at least one comparable has a net adjustment sign different from the others.

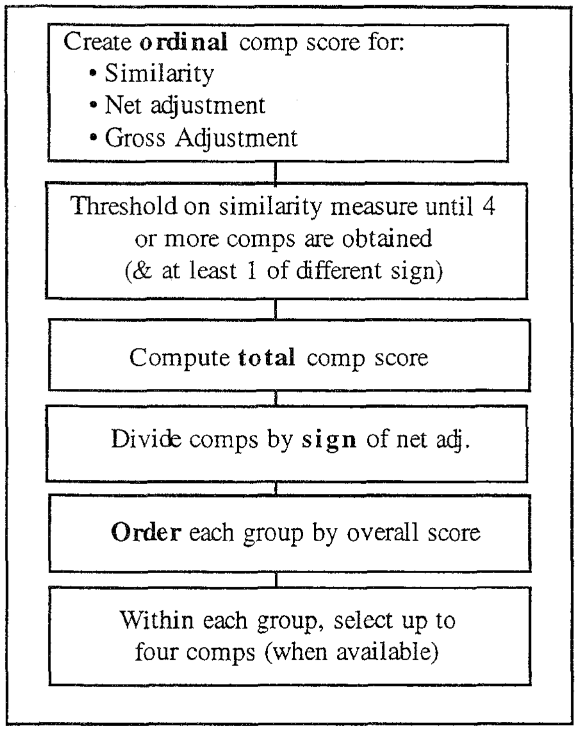

Figure 9: Comparable Final Selection

Comparable	Score Value	Score Rank	Net Adjust Value	N. A. Rank	Gross Adjust Value	G. A. Rank	Total Rank
113-012	0.95	1	1344	2	5924	4	7
306-018	0.88	2	3586	5	4186	1	8
093-011	0.78	3	5686	7	8191	7	17
305-006	0.67	4	6150	8	6160	6	18
685-046	0.64	5	3139	3	6099	5	13
847-984	0.58	6	-948	1	5670	3	10
873-005	0.53	7	-5261	6	9261	8	21
431-023	0.48	8	3546	4	4410	2	14
331-018	0.44	9	9310	9	11300	9	27

Figure 10: Comparable Selection Example

For the example in Figure 10, the comparables with the top six similarity scores would be included in the set. All other comparables are discarded. Of the comparables in the set we retain only four of each sign net adjust. The selected ones are the four with the lowest total rank. In the example from Figure 11, comparable number 305-006 would be discarded since there are four comparables with a positive net adjust and lower total rank. The five comparables selected form the final set of comparables.

4.6 Weighted Aggregation of Comparables Sales Price

After the best 4-8 comparables are found, their prices must be combined to produce the final estimate. Each comparable's contribution to this result is weighted by its similarity score. Figure 11 shows the calculation of a final estimate from adjusted values of the selected comparables and their similarity scores.

In addition to producing the final estimate of the value of the subject, the CBR estimator provides a confidence in the estimate and the comparables that justify the estimate.

Comparable	Adjusted Price	Score	Weighted Price
113-012	197000	0.95	187150
306-008	202000	0.88	177760
093-011	196500	0.78	153270
685-046	192000	0.64	122880
847-984	201000	0.58	116580
Total		3.83	757640
Final estimate = 757640/ 3.83			**199900**

Figure 11: Comparable Aggregation

5 Confidence Value Assessment

The users of PROFIT will make critical decisions based on the estimates generated. Therefore, we need to tell them when the system produces an accurate, reliable solution. We achieve this goal by attaching a confidence measure to each estimate. Ideally we would like to have subjects with the highest confidence exhibiting the lowest errors. At the same time we would like to assign high confidences to as many subjects as possible.

The confidence value is calculated from the following five quantitative characteristics of the case-based reasoning process:

- Number of cases found in the initial retrieval
- Average of the similarity values for the best four cases
- Typicality of problem with respect to the case-base (i.e. if the attributes of the subject fall within typical ranges for the subjects five digit zip code region)
- Span of adjusted sales prices of highest confidence solutions (i.e. the highest adjusted sale price minus the lowest adjusted sale price among the selected comparables)
- Distribution of adjusted sales prices of highest confidence solutions (i.e. average percentage deviation of the adjusted sales price of the comparables from the estimated value of the subject)

These characteristics are evaluated using the fuzzy membership functions illustrated in Figure 12. These functions map the numerical value of each parameter into a standard numerical confidence, which ranges from 0 to 1. These standardized confidence values are then aggregated into a final confidence value. Given the conjunctive nature of this aggregation, we decided to use the *minimum* of the standardized confidence values.

Figure 12a) shows that if two or less comparables are found then the standardized confidence for comparables found is 0. If between two and seven comparables are found, the confidence is ((n -2)*0.15)), i.e. the confidence increases 0.15 for each comparable over two to reach 0.75 when there are seven comparables. Between seven and twelve comparables, the confidence is ((n - 7)* 0.05) + 0.75), i.e. the confidence increases 0.05 for each comparable over seven and reaches 1.0 with twelve comparables. Since the aggregation method is minimum operator, a low confidence in any of the characteristics will cause a low confidence in the result regardless of other excellent confidences for the other characteristics. The other figures show similar membership functions for the other confidence measures.

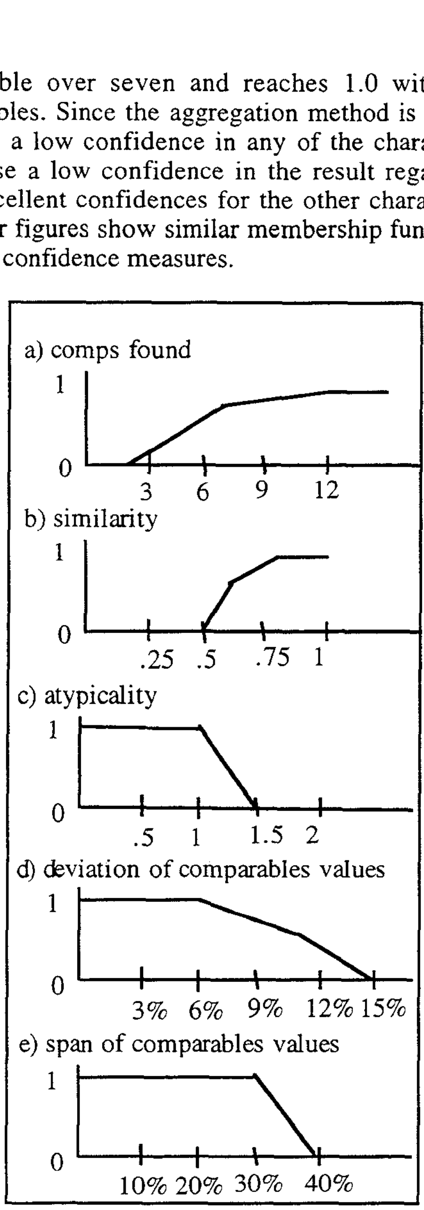

Figure 12: Membership functions for confidence

Figure 12b) shows that we have no confidence in an estimate whose similarity is lower than 0.5. Otherwise the confidence rises to 1 at a similarity of 0.8, and is 1 for anything over 0.8. Figure 12c) is used to determine our confidence based on the subject's atypicality. Atypicality is computed as a normalized deviation in the region's average lot size, living area, number of bedroom, bathrooms, etc. We have no confidence in subjects with atypicality greater than 1.5. Our confidence rises linearly as the atypicality decreases from 1.5 to 1.0 and is one when atypicality is less than 1.0. Figure 12d) shows our confidence for the average deviations in the values of the comparables. We have zero confidence in an estimate if the average comparable deviates from the estimated price by more than 15%. Finally, Figure 12 e) is used to determine our confidence based on the size of the span of the adjusted values of the comparables. If the span is greater than 40% of the value of the subject then we considered it too scattered and have no confidence in the estimate.

6 Statistical Analysis

To create the membership functions for PROFIT's five characteristics, which are illustrated in Figure 12, we ran our system on 7,293 properties from Contra Costa county in California. The predicted sales price of each property was calculated and compared with its actual sales price[3] to derive the estimate's error. The percentage error and its five confidence characteristics were calculated each subject. Figure 13 shows the values calculated for a random sample of ten of the 7,293 subjects. Each row is a different subject. The columns show the estimate error, the five characteristics calculated along with the estimate, and the confidence value obtained by taking the minimum of the evaluation of the membership functions of Figure 12 using the estimate's five characteristics.

Error	Comps Found	Simil.	Atyp.	Comps Dev.	Comps Span	Conf. Value
-9.8	3	0.63	1.42	2.02	6.32	**0.15**
-2	35	0.94	0.38	2.24	8.57	**1.00**
17.3	11	0.71	0.94	5.67	19	**0.70**
0.5	24	0.85	0.66	2.05	7.24	**1.00**
-1.6	14	0.95	0.29	2.89	9.33	**1.00**
5.2	15	0.90	0.73	3.24	12	**1.00**
5.2	12	0.74	0.17	4.5	18	**0.80**
3.1	19	0.74	0.81	2.83	8.11	**0.80**
-13.9	12	0.82	1.97	3.85	15	**0.00**
7.8	11	0.77	1.34	4.24	13	**0.32**

Figure 13: Sample of test run

Then we analyzed the conditional distributions of the estimate error, given each of its five confidence characteristics, and try to predict the error. For instance, Figure 14 shows that the estimate error (in percentage) decreases as the number of comparables found in the initial retrieval increases. Therefore, we can use this number as a filter to predict the expect expected error.

We used C4.5 [8] to create rules predicting the error from PROFIT's characteristics. Then we validated these rules via data visualization. Finally, the rules were manually transformed into the membership functions illustrated in Figure 12. The estimate's confidence value is the conjunctinve evaluation of all the rules.

At our customer's request, the confidence value generated by the rules was subdivided into three groupings (*good*, *fair*, and *poor*). The confidence measure should then produce the largest good set with the lowest

[3] Although the actual sale price was known, it was never used in the estimate's computation.

error. Of the 7,293 subjects, we could classify our confidence in 63% as *good*. The *good* set has a medium absolute error of **5.4%**, an error which is satisfactory for the intended application. Of the remaining subjects, 24% were classified as *fair*, and 13% as *poor*. The fair set has a medium error of **7.7%**, and the poor set has a median error of **11.8%**.

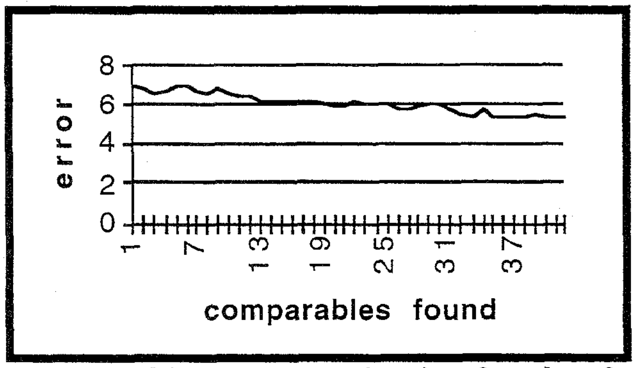

Figure 14: Median error as a function of number of comparables found

The PROFIT system can also be used in validation mode to support a sales price, or to determine the confidence value in a sales price. The sales price could either be given or supplied by another estimator. In this mode, PROFIT can find comparables to validate such a price. This feature is needed because appraisals are required to be supported by evidence of current market conditions. Using the actual sales price as a target, we were able to find comparables supporting a sales price with a median absolute difference from the actual sales price of **0.5%**. Clearly, when PROFIT is used in this mode it yields different results than when it is used to estimate the property's market value. In validation mode, the search for comparables is *biased* by the sales price. Since this price is also used to evaluate the results, PROFIT's error statistics show a large reduction in the median of the absolute error (from 5.4% to 0.5%). Note that the statistics shown in Figure 1, indicate the median of the error for Human Appraisers to be equal to 3%. These statistics were derived from the trade literature and include a large number of cases in which the sales price was available to the appraiser. On the other hand, not all sales price are necessarily originated from *bona-fide* transactions, so PROFIT results obtained from the *unbiased* search are actually more reliable.

7 Conclusions

We have shown the role of fuzzy logic in the development of PROFIT, a Case-Based Reasoning system for residential property valuation. Fuzzy logic is used in PROFIT's similarity computation, solution adaptation, and confidence value generation. As a result, the system has a very transparent process based on fuzzy logic's flexible, descriptive language used to translate current appraisers practices into PROFIT's retrieval criteria and adaptation rules. The system is also quite scaleable, as proven by the thousands of transactions used to test it. Finally, PROFIT outputs a confidence value associated with its estimate value. The confidence value, obtained from the conjunctive evaluation of a set of five soft constraints on the estimator's internal parameters, will guide the end user in determining whether to use PROFIT's estimate or to request an external valuation performed by a human appraiser.

PROFIT can also be used to validate a property value provided by an external source. In this mode the system identifies the best set of comparables to justify the given value and provides an associated confidence value.

Future work will include automatic case-base maintenance and update (the automatic determination of whether the selection or adaptation rules need to be changed, due to changing market conditions) and automatic generation of the new selection and adaptation rules. Finally, we would like to focus on representing and handling the uncertainty of cases caused by missing values, and integrating data sources of heterogeneous quality.

References

[1] Appraisal Institute, *Appraising Residential Properties, Part VI*, Chicago IL, 1994.

[2] Bonissone, P.P. and Ayub, S. "Similarity Measures for Case-based Reasoning Systems", Proc. of the Fourth Intl. Conf. on Information Processing and Management of Uncertainty (IPMU-92) in Knowledge-Based Systems, pp. 483-487, Palma, Spain.

[3] Breiman, Friedman, Olshen & Stone, *Classification and Regression Trees*, Monterey, CA: Wadsworth and Brooks, 1985

[4] Gonzalez, A. J. "A Case-Based Reasoning Approach to Real Estate Property Appraisal," *Expert Systems with Applications*, Vol. 4, pp. 229-246, 1992.

[5] Jang, J.S.R. ANFIS: Adaptive-Network-Based Fuzzy Inference System, *IEEE Trans. Systems, Man, Cybernetics*, 23(5/6):665-685, 1993

[6] Klir, G. and Folger, *Fuzzy Sets, Uncertainty, and Information*, Prentice Hall, 1988.

[7] Kolodner, J. *Case-based Reasoning*. Morgan Kaufmann, 1993.

[8] Quinlan, J.R. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, CA, 1993.

[9] Saaty, T.L. *The Analytic Hierarchy Process*, McGraw-Hill, 1980.