Case Age:
Selecting the Best Exemplars for Plausible Reasoning Using Distance in Time or Space

Mark H. Burstein
Bolt Beranek and Newman Inc.
10 Moulton St.
Cambridge, MA 02138
EMAIL: burstein@bbn.com

Abstract

The age of a case (in the CBR sense) is the amount of time that has elapsed between the time that the case originally occurred and the time of the current reasoning activity. People engaged in plausible reasoning tasks will, under appropriate circumstances, use the age of retrieved prior cases to filter and discard them, or to select among alternatives by their recency. This paper examines how the age of a case (and its spatial analog) are used by people in plausible and case-based reasoning tasks. I will argue that (1) The age of a retrieved case is an important factor in case relevance judgments for certain kinds of inferences. (2) When case age is relevant, more recent cases are usually, *but not always*, preferred to older ones (the "all other things being equal" caveat). Finally, I will argue that, somewhat surprisingly, (3) case age *cannot be used as an index into memory* given some commonly held assumptions about the nature of the retrieval process because it varies with the time of retrieval. This limits its use to post-retrieval processes, such as the filtering of already retrieved cases.

Introduction

Case age is the amount of time that has elapsed between the time that a case occurred and the current reasoning activity. This paper examines how the age of a case is used in case-based reasoning tasks. I will use anecdotal evidence to argue that:

- The age of a case is an important feature in similarity comparisons and judging the appropriateness of a case as a model in case-based reasoning tasks in many domains.
- When the recency of a case is an important guide to predictive power, people explicitly trade off age and other measures of similarity in choosing the most relevant case to use. In short, most recent is not always best.
- In most models of memory retrieval, case age (as opposed to a date, time, or epoch associated with the case) cannot be used as an index into case memory because the age of each case in memory is constantly changing. It is not a declarative feature available for creating indices at storage time. This restricts its use in plausible and Case-based reasoning tasks, as compared to most other features.

I will motivate the discussion with examples from three domains that I have examined directly or indirectly, plausible reasoning about recurring events like buses stopping at bus stops, and "expert" human case-based reasoning in the domains of real estate appraisal and labor contract negotiations

The Notion of Extent

The issue examined in this paper surfaced in our investigation of the role of proximity in time and space on human plausible reasoning (Burstein et al, 1992; Collins and Michalski, 1989; Collins, 1978). Collins (1978) identified a parameter he called "extent" which was particularly prevalent in temporal and spatial inferences. It is necessary because people have a notion of approximately how far things like rainstorms, parades, and large geographical features (e.g., plains, rivers) extend in space, and (where applicable) how long they extend in time. By maintaining these types of estimates, people can make judgments related to the so-called "frame-problem" (McCarthy and Hayes, 1981) or the "extended persistence problem" (Shoham, 1987).

Consider the following question, asked of a colleague:

Q: If you were away from BBN for a year, would you still expect your bus to come at 5:15?

A: Well, I don't know.. It comes fairly frequently around then, so I would expect it to come near that time. It might come then -- they don't change all of the scheduled buses that often.

At top level, the question goes to the extent of time that a prediction about this bus' arrival can be made directly based on its known past behavior. The answer shows three kinds of reasoning with uncertain information. One inference uses the age of last known case where the behavior occurred (e.g., whether he saw it come at 5:15 yesterday, or last week, or last year). The age of that information affects the certainty of his conclusion. Our model is that this age is compared to an internal expectation about the length of time that the pattern of the bus arriving at 5:15 can be expected to hold. A second line of reasoning is based directly on knowledge of the frequency of the phenomena *near that time* of day. This is a form of extent information, not unlike that used to

predict how long a rainstorm will last. Finally, there is knowledge of approximately how long an interval exists between changes to the bus schedules in general. This too is a piece of extent knowledge, related to a particular causal factor affecting the issue in question.

The case-based reasoning process, much like our model of human plausible reasoning (Collins and Michalski, 1989; Burstein et al. 1992) can be characterized as having the following steps:[1]

1. Retrieve some (finite number of) cases with similar features to the current situation.
2. From these, select the most appropriate one(s) by further comparisons with the current situation and with each other.
3. Use differences between the selected prior cases and the current situation on factors affecting the conclusion to be reached to adjust each mapped conclusion.
4. Combine (or select among) the conclusions reached with each case-driven inference.

In the remainder of this paper, I examine how the age of a case is used in reasoning by experts about past cases. Though I will not present specific evidence of it, most of the conclusions reached should apply equally to inferences that use the spatial distance between prior cases and the current one, when that measure is a predictor of similarity for an inference about the target.

The Various Forms of Age in Past Cases

There are several forms of "age" that can be associated with cases, or aspects of cases. These different forms behave differently with respect to their use in reasoning from prior cases to new ones. To clarify the discussion, I briefly distinguish all of the following:

- **Age of the object** of the case:
(e.g., Ford Mustangs start to have transmission problems when they are 4 yrs old.)
- **Vintage of the object** in the case: (e.g., '87 Chevys often have fuel pump problems.)
- **Age of the case:** (e.g., I replaced the fuel pump in my '87 Chevy last year.)

Each of these facts can be used in different kinds of case-based or plausible reasoning. In most plausible inferences, the similarity or relevance of a case is determined partially by knowledge of **dependencies**[2] (causal and/or statistical correlations between features). In domains like machine repair, the age of the object is potentially an important feature, because machine parts are sometimes known to wear out at predictable rates, as in the "mean time to failure" of the object. Similarly, several kinds of information about the date of manufacture (vintage) are commonly used as predictors of failure in automobiles and other machines. Model year is the most obvious one, but there are also the old stories about American cars produced on Mondays and Fridays where construction errors were above average.

These two time-related features are different than the age of the case, which is the primary focus here. Both the age and vintage of the artifact in the case can be used directly in retrieving relevant past cases because they can be computed and stored with the case. The age of the case itself is *the amount of time elapsed since the event occurred till the present.* . Note that if the reasoner has only secondary information about the event, then this will be different than memory storage time. It is most relevant in reasoning tasks where the thing being predicted or plausibly estimated is itself known to change over time. For example, in estimating the price of an item that is in many ways unique, the most recent case of a sale of a similar item is generally the most desirable prior case to base your conclusions on. Our examples from domain experts indicate that they form heuristic rules that explicitly deal with these tradeoffs.

The CBR Method in Real Estate Appraisal

The approach that real estate agents and professional appraisers use to estimate the sale value of a piece of property is a classic example of case-based reasoning. The following quote is from a textbook on the subject (Ventolo and Williams, 1987): *(italics added by this author)*

> "The direct market comparison or sales comparison approach to appraising uses the principle of substitution. The appraiser finds three to five (or more) properties that have been sold *recently* and are similar to the property that is the subject of the appraisal. The appraiser notes any dissimilar features, and makes an adjustment for each using the formula:
>
> **Sale price of comparable property**
> **+/- Adjustments**
> **= Indicated value of subject**
>
> Major types of adjustments include those made for physical (on-site) features, locational (off-site) influences, conditions of sale (e.g., financial terms) and *time from date of sale.*
>
> After going through this process for each of the comparable properties, the appraiser selects a value for the subject property that is the adjusted value of the comparable(s) most like the subject."

---

[1] See, e.g., Kolodner et al. (1985). A discussion of the retrieve-then-filter model can be found in Gentner and Forbus (1991). In the domains that I examine here, the retrieval process need not even be a memory-based process, as cases may be stored in paper or electronic form. The argument does not apply fully to techniques used for retrieval of cases stored in external media.

[2] *Dependencies* (Collins and Michalski, 1989), also called *determinations* by Russell (1986), are relationships between terms or attributes whose values are believed to be correlated in general. To use a dependency in a predictive inference, an example that is similar to the case in question must be found that provides a value on the correlated attribute. This value is then used as a predictor in the current case, under the assumption that "all else being equal, the attribute values should be the same." Dependencies can also be annotated in a number of ways, such as whether they are directly or inversely correlated. An example would be "Average annual temperature is inversely correlated with latitude." This allows for a wider variety of inference patterns.

The textbook talks about all of the features that must be taken into account in comparing properties. These include such things as the location, (location, location, ...); the date of sale; the size of the lot; the landscaping; type of construction; the age of the house (i.e., the *vintage*}; the kinds of improvements made; the style of home (e.g., the number of floors); numbers of rooms, bedrooms and baths; square feet of living and other space; exterior and interior condition, presence and size of a garage; and type of financing.

In selecting comparable properties, the book suggests that the location of the comparables should within the same neighborhood as the subject property, if at all possible (This is the analog of case age -- the spatial distance to the (object of) the prior case)

The **date of sale** of a property is very important in choosing whether or not to use it as comparable in an appraisal. The age of the house (the object of the case) is less important since most houses within a neighborhood will be built at approximately the same time, and an age difference of 5 years is not significant in most cases. When the date of sale of the comparable property occurred *more than six months prior* to the appraisal, adjustments must be made to the value of the comparable when estimating the value of the subject. In an interview with an experienced real estate agent who has done hundreds of appraisals, he said the following:

> Q: Would you ever use a comparable that sold more than a year ago?
>
> A: If it was an unusual property. The most recent ones are best. But in a situation where there weren't too many comparable properties that had sold recently .. a multi-family in an area of mostly single family homes for example.. yes. I'd have to do a lot of adjustment, though, for changes in the market.
>
> ... I would use a similar home that sold six months ago as a better comp than a not so similar one that sold last week.

In other words, an explicit trade-off is made between the degree of similarity determined using features of individual properties that affect their value, and similarity of value based on the age of the other cases (time of determination of sale prices of comparable properties). The latter is treated separately because it indicates the level of uncertainty in the appraiser's knowledge of the *current value* of the comparable property, which is ideally what he would use to estimate the subject's value. The general rule here is to use comparables whose value was fixed by a sale within the last six months, because they have a reasonably accurately known current value, and so no adjustment for the age of the case is necessary. Fortunately, no appraisal is ever based on only one comparable property. Unfortunately, since no two properties are quite identical, the appraiser must decide if the similarity of each comparable property outweighs the uncertainty of its current value, or if he or she can compensate for it. This in turn may depend on the overall volatility of the market. In a stable market, older sales may be more acceptable as comparables than in a volatile one.

The real estate agent I interviewed highlighted this point further in discussing exactly what the notion of the date of sale means:

> "In a rapidly changing market, you have to be careful estimating how old the sale was. A couple of years ago, when the market was changing fast, I did an appraisal where I didn't adjust for the age of a comp from six months before. The problem was that the purchase and sale agreement was signed 90 days before that. I valued it too high, and it didn't sell."

Thus, the real issue is the **time at which a purchaser and a seller agree on a price**, and hence determine its value. The definition of **age of case** in this specific reasoning task is the time elapsed since the last date when the sale price was fixed, which is the date of the purchase-and-sale agreement, not the date at which the sale was finalized, the recorded date of sale. The agent's error was in not understanding that the date of the setting of the price was the most accurate date to use in his measure of case age. Since the age of the comparable sale price was nine months, not six, the heuristic that an adjustment for "age of sale" should be used after six months was not applied when it should have been.

Case Age in the Adaptation of Comparable Values for Estimates

As the formula for adaptation of property values shown above suggests, case adaptation in real estate is a process by which adjustments are made to the sale price of a comparable property to compensate for differences from the subject. These adjustments are determined for most features, especially physical ones, essentially by an application of the principles of case-based reasoning to form adaptation rules by making plausible *generalizations* (Burstein et al, 1992) for local regions in time and space. I quote again from Ventolo and Williams (1987):

> "Ideally, if properties could be found that were exactly alike except for one category, the adjustment value of that category would be the difference in selling prices of the two properties. ...
>
> The appraiser can develop a Sales Price Adjustment Chart and use it to substantiate dollar amounts for adjusted values. Example:
>
> House A is very similar to house B, except that it has a two car attached garage and house B does not. Since the garage is the only significant difference between the two properties, its value is the difference between the selling price of $134,000 for A and the selling price of $125,500 for B. So the value of the garage is $8,500.
>
> ...
>
> *An appraiser will ordinarily need two or more instances of each variable as a check on the accuracy of an adjustment."*

This method is used with adjustments for the age of cases as well as with other properties:

"**Time adjustment.** Property H is identical in features to properties A, E, and J, except for the length of time since the date of sale. Property H was sold one year ago; properties A, E, and H were all sold within the past three months The adjustment value to be made in the case of a year-old sale is, therefore: $121,000 - $111,000 = $10,000 The $10,000 adjustment value could be expressed as a percentage of the total cost of the property, in this case 9% per year."

In my interview with a real estate agent, I asked about this method:

Q: Do you make price adjustments for age using dollar amounts or percentages?

A: Percentages, generally.

Q: Do you maintain a table of these percentages for all appraisals done around the same time, or by looking at changes in the values of similar properties each time you do an appraisal?

A: It depends. Similar kinds of properties in the same region, yes. You get a feel for how the market is going in an area. 20-20 hindsight is great, the problem is that you don't always know which way the market is going at the time.

From this and other things he said, it is clear that the cases he would consider making as comparables he could adjust for case age were also those from the same local region in space. The rate of inflation in property values was known to vary from town to town and neighborhood to neighborhood, and for different types of properties. He distrusted using comparables from different regions because of the uncertainty involved in mapping a sale price through too great a distance in both time and space at once, given knowledge of the dependencies relating price to those factors.

Case Age in Labor Contract Negotiations

To get a sense of the generality of the principles we observed, we examined a second domain that has been studied in depth using a case-based reasoning approach. Sycara (1987, 1988, 1991) developed a computer model of the role of an expert labor contract mediator. Her study was based on extensive interviews with one mediator, and the heuristic rules that she acquired from those interviews were reviewed and validated by two mediators.[3] These rules are summarized below:

- The most useful prior cases are ones negotiated within the same year. The best of those are cases were from negotiations for a competitor of the current company.
- Cases that were one to two years old and were for contracts of competitors to the current company are preferred to more recent cases for non-competitors
- Cases up to three years old can be used as precedents assuming that other similarity criteria are the same, and adjustments are made for time-sensitive features.
- Cases three to five years old can be used, but additional adjustments may be made based on a subsequent change in position of the company (in the case) in its industry. If the company has since improved its standing, then union economic demands in the new situation are adjusted upwards. If its position has declined, union demands are adjusted downward. If it has remained the same, no adjustment is needed.
- Cases five years old and older have no utility for economic concession negotiations, but may still be used for non-economic concessions, such as seniority provisions, work rules, and subcontracting provisions.

In this domain, the similarity of two cases is based first on whether the companies involved are direct competitors or not, and then on whether they are in the same industry. Two bus companies that serve some of the same areas are competitors. Two bus companies that don't compete are considered more similar to each other than to truck companies, though both are considered to be in the same "ground transport" industry. These are different again from sea transport businesses. A hierarchy of business types was used by Sycara to capture these distinctions, and to determine the relative similarity of various businesses. Another dimension of similarity is the set of job classifications that are being negotiated. Drivers of busses and trucks are considered similar in this dimension. If the industries are different, the similarity of their product lines and the economic conditions of the industries as a whole (boom, slump, or steady) are also considered.

Conclusions from Two Domains

The heuristic rules described above for real estate and labor mediation suggest several general principles are at work. First, because economic values (salary figures, real estate prices) are examples of time dependent features, plausible estimates of value developed by reasoning from historical cases where value was determined for similar objects can be made based on:

1. The most recent cases that are among the most similar on all other grounds. Estimates based on several prior cases are best.

2. Adjustments of each estimate (from a single other case) for the uncertainty in value of the object of the prior case in the time between the occurrence of the prior case and the present. Such adjustments add some additional uncertainty to the conclusion reached, since the change in value of the object in the prior case and the current one may not be the same.

[3]The two federal mediators who reviewed this work were of quite different backgrounds. One was a professor of economics at Georgia Tech. The other was a mediator that had prior experience with the Machinist's Union. Both concurred that the heuristics were accurate.

Each domain expert used a coarse model of the rate of change of these time-dependent values, that partitioned time into regions within which no adjustment for case age was required. These intervals represent a quantization of time for the dependency between case age and a dependent variable (e.g., salary rate or property value). A similar quantization of properties in space was also used (neighborhoods, towns) to rule cases out or rank them for similarity. Within a quantum, age is not a consideration. In real-estate, this is apparently six months under normal conditions. For labor contracts it is one year. Clearly these heuristic values for time periods are subject to change if general indicators like the rate of inflation fluctuate rapidly.

Finally, when prior cases are not equally similar to the current situation, the decision about which is the better model will have to trade-off the levels of similarity of each case to the target, and its recency. Explicit heuristics were developed by these experienced reasoners in their domains to handle the most common examples of this kind of trade-off reasoning. These heuristics explicitly rank cases meeting specific levels of similarity (e.g., competitor vs. non-competitor).above less similar but more recent cases, when measured in the time quanta established to approximate that domain's relevant age dependencies.

Case Age Cannot be a Memory Index

A final interesting point about the role of case age in case-based and plausible reasoning is that, unlike other features that can be associated with cases, the age of a case is not fixed at any particular time. Rather, **the age of a case depends on the time at which the case is retrieved**, which changes from moment to moment. Similarly, in situations where similarity may be correlated with spatial proximity, the relevance of a case depends on the proximity of its object to the object of the new situation that is the retrieval key. This is, in part, why we model all of the inferences described in this paper as happening *after* cases are retrieved. While it is quite common to store the *date of occurrence* of a case, in some form, as part of the case, this is not the same as the case age feature used in the reasoning described above, though it can be used to determine age after retrieval. This is consistent with the multi-stage model of retrieval proposed by Gentner and Forbus (1991).

Therefore, I would argue that age is not directly available as a means to select or filter which cases are actually retrieved. This violates a common assumption of many cognitive science researchers that the most similar cases can be retrieved in one step from memory. For that to be true, the memory system would have to *compute the distance in time* between the time of occurrence stored with the case and the present, and compare that value to the corresponding computed values for other potential cases in memory. This would place a heavy computational burden on the indexing and retrieval mechanism, and require that the time of occurrence of an event be treated much differently than other directly observed features in similarity comparisons.

It has been argued that connectionist models can, in fact, exhibit this behavior, if there is decay of case memories over time. There are two potential problems with this suggestion: First, it only works for direct experiences. If the case in memory was acquired second hand, then there is no correlation between the original time of occurrence and the time of memory encoding. Thus, the amount of "decay" would have to be related to the *reported* age of the case, not the time of its acquisition. Second, the degree of similarity of two cases due to their age is very domain dependent. This is clear from the differing cutoff heuristics people use for how similar ages must be in order to ignore age as a factor. This too would have to be factored into any connectionist model that attempted to directly compare cases by age during retrieval.

These arguments do not rule out the use of a number of memory retrieval strategies to restrict the set of cases considered, such as by retrieving only cases identified as having occurred in specific **eras** (Kolodner, 1983) (e.g., "any relevant cases in the last year? the previous year?..."). Such strategies apply equally when the cases are not actually stored in memory, as is generally the case in fields where precedents are important. Indeed, the experts interviewed suggest that they thought it important to see cases that weren't the most recent ones, in order to make explicit trade-offs in their choice of analogs.

References

Burstein, M. H., Collins, A. and Baker, M. (1992). Plausible Generalization: Extending a Model of Human Plausible Reasoning. *Journal of the Learning Sciences, 1(3), 319-359.*

Collins, A. (1978). Fragments of a theory of human plausible reasoning. In D. Waltz (Ed.), *Theoretical Issues in Natural Language Processing II* (pp.194-201). Urbana, IL: University of Illinois.

Collins, A. and Michalski, R. S. (1989) The logic of plausible reasoning: A core theory. *Cognitive Science*, 13, 1-49.

Gentner, D. and Forbus, K. (1991) MAC/FAC A Model of Similarity-based Reasoning. In *Proceedings of the Thirteenth Annual Cognitive Science Conference*, U.Chicago. (pp. 504-9). Hillsdale N.J.:L. Erlbaum Associates

Kolodner, J.L. (1983) Maintaining Organization in a Dynamic Long-Term Memory. *Cognitive Science*, 7(4).

Kolodner, J.L., Simpson, R.L. and Sycara-Cyranski, K. (1985) A Process Model of Case-Based Reasoning in Problem Solving. In *Proceedings of the Ninth IJCAI.* (pp. 284-90).Los Angeles: Morgan Kaufman.

McCarthy, J.M. and Hayes, P.J. (1981) Some Philosophical Problems from the Standpoint of Artificial Intelligence. In *Readings in Artificial Intelligence,* (pp. 431-450) Palo Alto:.Tioga Publishing Co.

Russell, S. J. (1986) *Analogical and Inductive Inference.* Doctoral Dissertation. Palo Alto: Stanford University, Dept. of Computer Science.

Shoham, Yoav. (1987) What is the frame problem? In *Reasoning about Actions and Plans.* (pp. 83-98). Los Altos, CA: Morgan Kaufman.

Sycara, Katia. (1987) *Resolving Adversarial Conflicts: An Approach Integrating Case-Based and Analytic Methods.* Doctoral Dissertation, Atlanta: Georgia Institute of Technology, School of Information and Computer Science.

Sycara, Katia. (1988) Patching Up Old Plans. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society,* Montreal, Canada

Sycara, Katia. (1991) Problem Restructuring in Negotiation. *Management Science,* 37:10, 1248-1268.

Ventolo Jr., William L. and Williams, Martha R. (1987) *Fundamentals of Real Estate Appraisal (4th Edition).* Real Estate Education Company.

Acknowledgments

I would like to thank Katia Sycara for the examples she contributed, and she and Bill Ferguson for discussions on the issues presented.