Analogical Estimation

Praveen Paritosh (paritosh@northwestern.edu)
Qualitative Reasoning Group, Department of Computer Science,
Northwestern University, 1890 Maple Ave,
Evanston, IL 60201 USA

Abstract

Analogy lets us make inferences from a better known example on to a less known one. Analogical estimation is a specific kind of analogical inference, namely, inferring the quantitative value of an unknown based on a known value from a similar example. We present *Analogical Estimator*, a computational model of analogical estimation. A key claim of this work is that we don't need a separate mechanism to capture the role of quantitative values in similarity. We need better representations, specifically, symbolic representations of quantity that capture the essential qualitative distinctions for that quantity.

1 Introduction

Consider the following examples:

- How much will I spend on this year's conference trip?
- What is the engine power of the Honda Civic?
- What is the population of Austria?

One of the ways to answer these questions is using a similar example to get a rough estimate. If there is a conference trip in my memory, which is similar enough to this year's, I might say that as a rough estimate we will spend the same amount of money. If I knew Austria was similar enough to Switzerland [Brown and Siegler, 2001], with a population of 7.5 million, I could guess the same value for Austria[1]. These are inferences of the value of a quantity from a similar example. This is a powerful method to get a rough estimate of an unknown quantity, as long as I can retrieve a good analog. Let's look at the conference trip example. This year the Cognitive Science conference is in Italy, last year it was in Chicago. Even though last year's Cognitive Science conference might have a high similarity to this years, it is not a good example to analogize from in the context of estimating expenses since I live in Chicago, which meant I didn't have to spend any money on travel or accommodation. A good example from my experience for estimating expenses is the last time I went to Italy, which was for a different conference. In this situation, we would want the analogical estimation to prefer that over last year's Cognitive Science conference.

Doing this kind of inference poses the following questions –

1. *Similarity*: How to take quantitative values into account in both retrieval and matching processes?
2. *Inference:* What sanctions and supports a quantitative inference?

In this paper, we present *Analogical Estimator*, a computational model of making quantitative estimates by analogy. The Analogical Estimator is based on cognitively plausible answers to the two questions posed above. A key claim of this work is that we don't need a separate mechanism to capture the role of quantitative values in similarity. We need better representations, specifically, symbolic representations of quantity that capture the essential qualitative distinctions for that quantity. Bringing together evidence in linguistics and psychology, we argue that our representations must make two kinds of distinctions – *dimensional*, those that denote changes of quantity, e.g., large and small; and *structural*, those that denote changes of quality, e.g. boiling point and poverty line. The Analogical Estimator offers a functional argument for the efficacy of these representations. We compare the accuracy of analogical estimates between representations enriched with the symbolic representations above, and without.

There are two substantial bodies of research – quantitative inference, and similarity, that are often not much related to each other. In the next section, we briefly review some of the relevant work in these fields. Section 3 presents a sketch of our approach, and section 4 talks about the model in more detail. In section 5 we present some results and discussion. We conclude with future work in section 6.

2 Background

This section reviews relevant research in the computational modeling and cognitive psychology of phenomena related to answering quantitative questions like "What is the population of Austria?" and "How much will I spend on this year's conference trip?" There are at least three different construals of the reasoning that could happen in answering these questions, leading to different bodies of research:

2.1 Problem solving

This approach to answering such questions uses heuristics and strategies [Larkin *et al*, 1988; Collins and Michalski, 1989]. We believe such problem solving is a key aspect of quantitative estimation and have developed a problem solver [Paritosh and Forbus, 2004] that uses strategies to decompose the problem into easier ones. But eventually, such problem solving has to bottom out with answers that

[1] A very interesting issue here is *adaptation:* adjusting the estimate, e.g., saying that Austria might be slightly larger than Switzerland, so I will guess 8 million. The Analogical Estimator does not do this as of yet.

are all known. A question that's left unanswered by a purely problem solving based approach is how do we produce reasonable guesses when we cannot transform the problem into subproblems whose answers are exactly known.

2.2 Numerical reasoning

This approach takes the perspective that the key reasoning involves numerical facts. A robust finding is the anchoring bias [Brown, 1953; Tversky and Kahneman, 1974; Kahneman, 1992]. One demonstration of the anchoring bias involves the subject making a comparison with an incidental number, called the anchor. Later on, when subjects are asked to come up with a quantitative estimate, then their answers are biased towards the anchor they were initially given. For example, participants were asked to compare the percentage of African nations in the UN as being as higher or lower than an arbitrary number (25% or 65%). Following this, they were asked to estimate the percentage of African nations in the UN. The mean estimates for the subjects who received the high anchor was 45% compared to 25% for the low anchor. Anchoring effects have been found with completely unrelated anchors, e.g., the sum of the last three digits of phone number added to 400 while estimating the year in which Attila the Hun was defeated [Russo and Shoemaker, 1989]. Anchoring effects have been found with both domain experts and novices, e.g., real estate agents and students estimating an appraisal value for a house after touring through it [Northcraft and Neale, 1987].

There is a growing body of evidence [Mussweiler and Strack, 2001; Chapman and Johnson, 1999] supporting that anchoring is not a purely numeric phenomenon, but has semantic underpinnings. Mussweiler and Strack's (2000) Selective Accessibility model of anchoring suggests that the anchor causes increased accessibility of anchor-consistent knowledge. For example, with the high (65%) anchor in the Africa example, facts like "Africa is a large continent" and "There are more African countries than I keep in mind" are retrieved [Trope and Liberman, 1996]. The final numeric estimate is generated based on the easily accessible knowledge [Higgins, 1996], so their estimate is heavily influenced by anchor-consistent knowledge.

Brown and Siegler (1993) used a different paradigm, which perhaps looks more like real-life quantitative estimation. Their typical experiment has three phases – first, participants are presented a set of items and asked to estimate the value of a particular quantity (e.g., populations of countries). Next, they learn the actual value of a subset of items (called *seed items*). Finally, the subjects re-estimate the values of items in the initial set. They found improved estimation as a result of seeding [Brown and Siegler 1993, 1996; Friedman and Brown, 2000; Walbaum, 1997]. By manipulating the seed items, they found that people access two independent sources of knowledge while generating estimates – (1) *Metric knowledge:* information about the numerical properties of the quantity, and (2) *Mapping knowledge:* non-numerical information about the domain which could be used to order items relative to one another with respect to that quantity. Brown and Siegler (2001) have shown that seeds behave differently than anchors. Seeds provide both metric and mapping knowledge by providing feedback ("small European countries have fewer people than I would have guessed"). In contrast to anchors, seeds can push estimates of target items away from itself.

These data suggest that quantitative estimation is not a purely numeric task. Non-numeric knowledge is used to construct estimates. Furthermore, the integration of the numeric and semantic has not been addressed much.

2.3 Similarity

There are models in case based reasoning [Ashley, 1990; Leake, 1996; Ram and Santamaria, 1997] that use numeric information, but they employ *ad hoc* similarity metrics such as Euclidean distance that are not psychologically grounded. In the domain of estimation of length of software projects, ANGEL [Shepperd and Schofield, 1996] makes an estimate by retrieving a similar project. The similarity is defined by a numeric error metric that is minimized. However successful some of these applications might be, we don't think they tell us much about how numeric quantities are implicated in similarity judgments.

There is converging psychological evidence for structured models of retrieval, similarity and generalization. The structure-mapping engine (SME) [Falkenhainer et al, 1989] is a computational model of structure-mapping theory [Gentner, 1983]. Given two structured propositional representations as inputs, the base (about which we know more) and a target, SME computes a mapping (or a handful of them). MAC/FAC [Forbus et al, 1995] is a model of similarity-based retrieval, that uses a computationally cheap, structure-less filter before doing structural matching. It uses a secondary representation, the content vector, which summarizes the relative frequency of predicates occurring in the structured representation. The dot product of content vectors for two structured representations provides a rough estimate of their structural match. SEQL [Kuehne et al, 2000] provides a framework for making generalizations based on computing progressive structural overlaps of multiple exemplars.

One limitation of these models – and of other models of analogical processing (e.g., ACME [Holyoak and Thagard, 1989, LISA [Hummel and Holyoak, 1997], ABSURDIST [Goldstone and Rogosky, 2002]) – is that they do not handle numerical properties well. In most of these models, numbers are treated like symbols, so 99 and 100 are as similar/different as 99 and 10000, when treated as symbols, they are both non-identical symbols, but numerically, the differences in magnitude are quite different. As a consequence, we have following limitations in the retrieval, matching and generalization processes:

Retrieval: Just as Red occurring in the probe might remind me of other red objects, a bird with wing-surface-area of 0.272 sq.m. (that is the Great black-bucked gull, a large bird) should remind me of other large birds. This will not happen in the current model, unless we abstract the numeric representation of wing-surface-area to a symbol, say, Large.

Similarity: A model of similarity must be sensitive to quantity. For example, in current matchers, two cars which are identical in all dimensions have the same similarity as two that differ in some dimensions, if other aspects of their representations are identical.

Generalization: A key part of learning a new domain is acquiring the sense of quantity for different quantities. E.g., from a trip to the zoo, a kid probably has learnt something about sizes of animals.

3 The approach

The main hypothesis driving our approach is that we can model the role of quantities in similarity if we have structured representation of quantitative knowledge, namely, symbolic qualitative representations that capture the important distinctions. Consider two descriptions being compared that have a bunch of quantitative attributes. Two questions arise in computing the contribution to the overall similarity between the two descriptions:

3.1 How to compute similarity along a dimension?

Most models that use numeric information compute dissimilarity, which is usually a difference of the two values.

We claim that an essential part of computing similarity of two quantitative values along a dimension are – (1) comparing them to important reference points on that dimension, and (2) comparing their equivalent linguistic representations, e.g., are they low, typical, or high values. The psychological reality of reference points on the scale of quantity has been shown in various domains. Rosch (1975) argued for the special status of such "cognitive reference points" by showing an asymmetry – namely that a non-reference stimulus is judged closer to a reference stimulus (e.g., the color off-red to basic-red) than otherwise, while such relationship between two non-reference stimuli is symmetric. Existence of landmarks to organize spatial knowledge of the environment and similar asymmetries have been found [Holyoak and Mah, 1984 among others]. Other relevant psychological studies that support the existence of reference points come from categorical perception [Harnad, 1987] and sensitivity to landmarks [Cech and Shoben, 1985]. Our knowledge of how the quantity varies (its distribution), and linguistic labels like `Large` and `Small`, are but a compressed record of large number of such comparisons. The semantic congruity effect [Banks and Flora, 1977] is the fact that we are better and faster at judging the larger of two large things than the smaller of two large things. Part of the account from experiments involving adults learning novel dimension words, by Ryalls and Smith (2000) is the fact that in usage, we make statements like "X is larger than Y" more often than "Y is smaller than X", if X and Y are both on the large end of the scale.

3.2 How to combine similarities along different dimensions?

In existing computational models, this is usually accomplished by a Minkowski metric combination of differences along various dimensions. Sometimes, domain specific weights might be used to weigh different dimensions in a variable manner [Nosofsky, 1991].

The causal status hypothesis [Ahn *et al*, 2000] suggests that a feature is more important and should be weighed higher if it is deeper in the causal chain of relationships to other features. Feature centrality [Hadjichristidis *et al*, 2004] claims that features on whom many other features depend will be more likely to be carried over to a target concept. The structure mapping theory [Gentner, 1983] already has the most general version of this idea. The contribution of two similar elements is weighed by the structural support for the match of the elements: if the two elements are involved in deeper relational structure, they are judged more similar. For example, two basketball players being of similar height leads to a stronger contribution of overall similarity then two professors being of similar height; since height is involved in larger relational structure of the domain of basketball players than professors.

So, our answer here is that the combination of similarities from different dimensions is automatically a given in the structure mapping framework. The structural support also gives us a solid reason to believe that an inference can be made in the target based on the fact in the base case. In the next section we describe what these essential qualitative representations are, and CARVE, a computational model to generate the representations automatically from a set of examples.

4 Representing Quantity

Our knowledge about quantities is of various kinds – we understand that there are `Expensive` and `Cheap` things, that Canada is `larger` (in area) than the USA, that basketball players are usually `tall`, that the `boiling point` of water is `100 degrees Celsius`. A key part of such knowledge seems to be a *symbolization* of the space of values that a quantity can take. A symbolization identifies and names intervals and points in the space of values of a quantity. We argue that our representations must make two kinds of distinctions – *dimensional*, those that denote changes of quantity, e.g., large and small; and *structural*, those that denote changes of quality, e.g. boiling point and poverty line.

4.1 Constraints on cognitive representations of quantities

A representation of quantity allows us to make certain distinctions – numbers allow us to make too many, and dividing the range of values into two equal sized parts doesn't necessarily provide useful distinctions. Representations do not arise in vacuum. Our representational framework must be capable of capturing the interesting ways in which a quantity varies in real-world instances of it. Below are two different kinds of constraints on values a quantity can take –

Distributional Constraints

Most quantities have a range (a minimum and a maximum) and a distribution that determines how often a specific value shows up. For example, the height of adult men might be between 4 and 10 ft, with most being around 5-6.5ft. We can usually talk about the `low, medium, high` for many quantities, which seems to be a qualitative summary of the distributional information. There is psychological evidence that establishes that we *can* and *do* accumulate distributions of quantities [refer to Malmi and Samson, 1983; Fried and Holyoak, 1984; Kraus *et al*, 1993; among others, for more].

Structural Constraints

Quantities are constrained by what values *other* quantities in the system take, its relationship with those other quantities, via its relationships with them[2]. For instance, for all internal combustion engines – as the engine mass increases, the Brake Horse Power (BHP), Bore (diameter), Displacement (volume) increases, and the RPM decreases. These constraints represent the underlying mechanism, or causal model of the object. The *structural constraints* on quantities reflect a fundamental fact about the way things are in the world. Things in the world come in *packages* or *bundles.* For example, a "muscle car" has a powerful engine, is expensive, is designed for style and fun rather than safety or practical driving. In psychological literature, a similar notion is expressed by *attribute co-variation* or *feature correlation* [Malt and Smith, 1984; Kersten and Billman, 1992 and McRae, 1992]. But there's much more than that – these are not merely bundles of correlated attributes, but are *structural bundles.* The entities, and quantities associated with them, tied by relations and higher order relations constraining them, give rise to the structure[3] therein.

These two ecological constraints point us to the two different kinds of information about quantities, which must be parts of our representations –

1. Distributional information about how the quantity varies.
2. Its role in and relationship to the underlying structure/mechanism, and the points at which there are changes in underlying structure.

4.2 Proposed Representation

There are two kinds of distinctions that our representation of quantity must make –

1. *Dimensional partitions:* Symbols like `Large` and `Small`, which arise from distributional information about how that quantity varies.
2. *Structural Partitions:* Symbols like `Boiling Point` and `Poverty Line`, that denote changes of *quality*, usually changes in the underlying causal story and many other aspects of the objects in concern.

[2] Comic books, mythology, and fantasy, for example, have the freedom to relax this constraint – a character can be arbitrarily strong, large, small or be able to fly, even though the physical design of the character might not be able to support it.

[3] The structure of relationships is an even more general notion than causality, spatial arrangement, connectivity.

4.3 CARVE: Modeling learning about quantities

Dimensional Partitioning

CARVE takes as input a set of cases represented in predicate calculus. For each quantity, it extracts all the numeric values for it in the input cases. Given these values, the job of the dimensional partitioning step is to find three partitions, corresponding to `Low`, `Medium` and `High` ranges of the values that the quantity takes.

These partitions are currently generated using a k-means clustering algorithm. It is possible to plug in different heuristics that partitions the values into ranges of values.

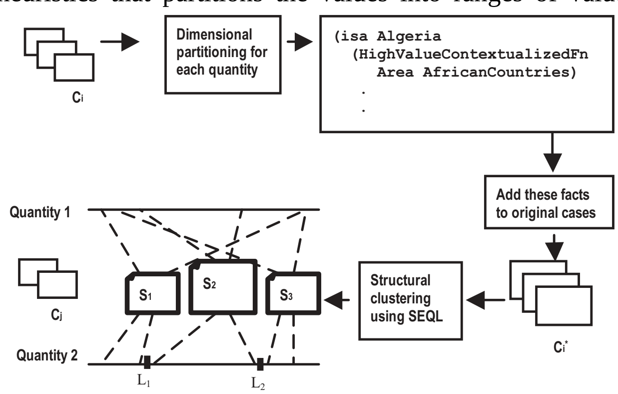

Figure 2. A schematic overview of how CARVE computes dimensional and structural partitions. The input is examples represented as cases C_i. Each case consists of a set of assertions in predicate calculus. The first step of dimensional partitioning leads to creation of new assertions, which are added to the original cases. In the structural partitioning step, SEQL is used to find the structural clusters, S_i, which are projected on to quantity dimensions to find the structural partitions. In the above example, **Quantity2** has structural partitions at L_1 and L_2, but **Quantity1** has no such partitions.

Heuristics based on central tendency and percentiles do not work for zipf-like distributions which we see in many of the quantities. For such distributions, means and variances are not intuitively meaningful at all. The k-means clustering algorithm fits with what people did in a pilot experiment we conducted [Paritosh, 2004]. On an average across subjects, the dimensional partitions computed by CARVE agreed with people 74% of the times (*sd*=27). More empirical data is needed to conclude what set of heuristics people use to make these partitions, and when they work.

For each fact about the value of a quantity, we then add a `High/ Medium/ LowContextualizedValueFn` to the case depending upon which range that numeric value fell in. These facts are used in the next step.

Structural partitioning

SEQL [Kuehne *et al*, 2000] provides a framework for making generalizations based on computing progressive structural overlaps of multiple exemplars. The goal of structural partitioning is to find the structural clusters in the cases and project these clusters on to various quantity dimensions. The cases produced at the end of the dimensional partitioning step are given as input to SEQL. In figure 2, we see the output of SEQL as three generalizations S_1, S_2 and S_3 and some leftover cases that did not fit any of those. Let's consider two quantities $Quantity_1$ and $Quantity_2$. The projection of a cluster on a quantity is the range of values for that quantity in the cluster. For $Quantity_1$, we see that the projections from all the three generalizations overlap. On the other hand, the projections of the generalization on $Quantity_2$ are non-overlapping. L_1 and L_2 denote the boundaries for these ranges. Notice the predictive power of knowing that for a specific case the value of $Quantity_2$ is less than L_1. We not only know about the quantity value, but about the generalization to which the case belongs, and so can predict other causal properties of it. For instance, when you know that a country is a developing country, there are rich causal predictions you can make.

5 Estimating Basketball Stats

To illustrate the above ideas, we report results from the domain of estimating stats (e.g., Points per game, Assists per game, height, etc.) of basketball players. For every player, a host of stats are available[4] and kept track of by fans and experts. There are some interesting causal relationships between various quantities, e.g., being tall helps to rebound and block. Players that are closer to the basket are typically taller. People further away from the basket have to be good shooters, and usually make more three-pointers. The point guard, whose key function is to dribble and pass is usually smaller, and has high assists per game. We selected fifteen players such that they were reasonably different, three from each of the five positions on the field. We built a case library in which each basketball player was represented as a case. We will refer to this case library as the raw case library. This initial library had eleven facts like the following for each player:

```
(seasonThreePointsPercent
     JasonKidd 0.404)
(qprop seasonThreePointPercent
     seasonFreeThrowPercent
     BasketballPlayers)
```

We ran these cases through CARVE to generate dimensional partitions for all the quantitative attributes, e.g.,

```
(isa JasonKidd
   (HighValueContextualizedFn
        seasonThreePointsPercent
        BasketballPlayers))
```

`High/Medium/LowValueContextualizedFn` are functions that take two arguments – a quantity and a context argument and return a collection of objects. So in the above example `HighValueContextualizedFn` denotes the collection of basketball players with a high season points per game, and the isa statement says that Jason Kidd is an instance of that collection. The `LowValueContextualizedFn` similarly lets us represent the negative end, for instance small and cheap. We create a new, enriched case library that consists of all the facts in the raw case library plus these newly generated assertions to the respective cases. In this experiment, no structural partitions are found, which means that we need an even larger case library with richer causal structure to get to structural partitions.

We then generate estimates using both raw and the enriched case library that are compared in Table 1. The setup is very simple. Suppose we want to estimate the height of Jason Kidd. We remove all statements regarding his height from his case, and then use that as a probe with both the enriched and the raw case libraries to remove the most similar analog. We then use the height of the retrieved analog as an estimate for Jason Kidd's height. The error of this estimate is then simply the difference between the estimate and the real value of height. The first promising thing is the overall error percent across all the estimates is half (16.9% as opposed to 34.5%) when we use the enriched cases.

Dimension	Enriched mean % error	Enriched error stdev	Raw mean % error	Raw error stdev
Height	1.4	1.7	6.1	3.1
Assists	17.7	39.2	68.4	82.1
Free throws	41.5	9.8	9.1	6.6
Points	15.9	12.9	28.4	29.8
Rebounds	14.4	17.7	52.3	37.5
Three point %	20.9	34.0	46.2	24.4
All	16.9	25.7	34.5	45.8

Table 1. We look at percentage errors across all estimates along different dimensions. The first two columns are errors and standard deviations for estimates using the enriched cases and the latter two using raw cases.

6 Conclusions

We presented a theory of analogical estimation that integrates quantity into similarity and analogy, by using cognitively plausible symbolic representations of quantity. We discussed these representations of quantity and a computational model, CARVE, which models learning about quantities. We then reported our results with estimating basketball stats using the above ideas. We believe that this paradigm of analogical estimation is an effective way to study the role of quantity in our judgments. The results are promising, and some of the next steps in this line of work are:

1. Generating confidence values for estimates based on structural support.

[4] http://www.nba.com

2. Enlarging our case library with richer representations and more examples.
3. Adapting the retrieved estimate to get an even closer estimate.

Acknowledgments

This research is supported by the Computer Science Division of the Office of Naval Research. The author would like to thank Ken Forbus and Dedre Gentner for discussions on these ideas, and Matt Klenk for help with the basketball examples.

References

Ahn, W., Kim, N.S., Lassaline, M.E., and Dennis, M.J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology,* 41, 361-416 (2000).

Ashley, K.D. (1990). Modeling Legal Argument, MIT Press, MA.

Brown, D. R. (1953). Stimulus-similarity and anchoring of subjective scales, *American Journal of Psychology,* 66, 199-214.

Brown, N.R and Siegler, R.S. (2001). Seeds aren't anchors. *Memory and Cognition,* 29(3), 405-412.

Brown, N.R and Siegler, R.S. (1993). Metrics and Mappings: A framework for understanding real-world quantitative estimation. *Psychological review.* 1993. 100(3). 511-534.

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

Cech, C. G. and Shoben, E. J. (1985). Context Effects in Symbolic Magnitude Comparisons. *Journal of Experimental Psychology: Learning, Memory and Cognition,* **11**, 299-315.

Chapman, G.B. and Johnson, E.J. (1999), Anchoring, activation and the construction of values. *Organizatinal Behavior and Human Decision Processes*, 79(2), 115-153.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence,* **41**, 1-63.

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science,* **19**(2), 141-205.

Friedman, A. and Brown, N.R. (2000). Reasoning about geography. *Journal of experimental psychology: General,* 129, 193-219.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7,* 155-170.

Goldstone, R. L. and Rogosky, B. J., (2002). Using relations within conceptual systems to translate across conceptual systems, Cognition, 84, 295-320.

Hadjichristidis, C. Sloman, S., Rosemary, S. and Over, D. (2004). Feature centrality and property induction. *Cognitive Science,* 28 (2004), 45-74.

Harnad, S. (1987). *Categorical perception.* Cambridge: Cambridge University Press.

Holoyak, K. J., and Mah, W. A. (1984). Cognitive Reference Points in Judgments of Symbolic Magnitude. *Cognitive Psychology*, **14**, 328-352.

Holyoak, K. J. and Thagard, P. R. (1989). Analogical Mapping by Constraint Satisfaction, Cognitive Science, 13, 295-355.

Hummel, J.E and Holyoak, K. J. (1997). Distributed representations of structure: a theory of analogical access and mapping, Psychological Review, 104, 427-466.

Higgins, E.T. (1996). Knowledge Activation: Accessibility, applicability, and salience. In E.T. Higgins and A.W.Kruglanski (Eds.), *Social Psychology: Handbook of basic principles* (pp239-270). New York: The Guilford Press.

Kahneman, D. (1992). Reference points, anchors, norms, and mixed feelings. *Organizational behavior and Human Decision Processes,* 51, 296-312.

Kuehne, S., Forbus, K., Gentner, D. and Quinn, B.(2000) SEQL: Category learning as progressive abstraction using structure mapping. *Proceedings of CogSci 2000*.

Larkin, J. H., Frederick R., Carbonell, J. and Gugliotta, A. (1988). FERMI: A Flexible Expert Reasoner with Multi-Domain Inferencing, *Cognitive Science*, **12**(1), 101-138

Mussweiler, T. and Strack, F. (2001). The semantics of anchoring. *Organizational Behavior and Human Decision Processes,* 86(2), 234-255.

Mussweiler, T. and Strack, F. (2000). The use of category and exemplar knowledge in anchoring tasks. *Journal of Personality and Social Psychology*, 78, 1038-1052.

Nosofsky, R. M. (1991). Stimulus bias, asymmetric similarity, and classification. Cognitive Psychology, 23, 94-140.

Northcraft, G.B. and Neale, M.A. (1987). Experts, amateurs, and real estate: An anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behavior and Human Decision Processes,* 39, 84-97.

Paritosh, P.K. and Forbus, K.D. (2004). Using Strategies and AND/OR Decomposition for Back of the Envelope Reasoning. In *Proceedings of the 18th International Workshop on Qualitative Reasoning,* Evanston.

Paritosh, P.K. (2004). Symbolizing Quantity. In Proceedings of the 26th Cognitive Science Conference, Chicago.

Ram, A. and Santamaria, J.C. (1997). Continuous case-based reasoning. *Artificial Intelligence,* **90**, 25-77

Russo, J.E. and Shoemaker, P.J.H. (1989). *Decision Traps.* New York: Simon and Schuster.

Rosch, E. (1975). Cognitive Reference Points. *Cognitive Psychology*, **7**, 532-547.

Shepperd, M.J. and Schofield, C. (1996). "Effort estimation by analogy: a case study," presented at 7th European Software Control and Metrics Conference, Wilmslow, UK.

Trope, Y. and Liberman, A. (1996). Social hypotheis testing: Cognitive and motivational factors. In E.T. Higgins and A.W.Kruglanski (Eds.), *Social Psychology: Handbook of basic principles* (pp239-270). New York: The Guilford Press.

Tversky, A., and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases, *Science*, **185**, pp 1124-1131.

Walbaum, S.D. (1997). Seeding the fast food knowledge base: Long term affects. Poster presented at the 38th Annual Meeting of the Psychonomic Society.