

# Cognitive Processes in Quantitative Estimation: Analogical Anchors and Causal Adjustment

Praveen K. Paritosh (paritosh@northwestern.edu)

Matthew E. Klenk (m-klenk@northwestern.edu)

Qualitative Reasoning Group, Northwestern University, Evanston IL 60208

## Abstract

Quantities are ubiquitous and an important part of our understanding of the world. How do people solve quantitative estimation problems? A large body of psychological research addressing this question is based on the anchoring and adjustment paradigm (Tversky and Kahneman, 1974). Based on an analysis of verbal protocols of experts engaging in estimation tasks, we claim that similarity and analogy play important roles in quantitative estimation. A similar example for which the answer is known provides an *analogical anchor*, and comparison provides the grist for making *causal adjustments*. We present KNACK, a computational model of analogical estimation. Our theoretical analysis of causal adjustments suggests that they might not always be insufficient, as they are based on a discrepancy between assumed strength of causal relationships and that exists in the world (Kareev et al, 1997).

## 1 Introduction

Making rough quantitative estimates is a key component of commonsense reasoning about everyday situations. Let's look at some examples of quantitative estimation problems:

- What is reasonable price for a 2001 Ford Focus hatchback with 41,000 miles on it?
- How much does a two bedroom apartment in the Rogers Park neighborhood of Chicago cost?
- What is the average points per game scored in the current season by Jason Kidd?
- What is the freezing point of Vodka?

There are many different factors involved in solving these types of questions. First, there is domain specific knowledge in form of rules and/or examples, e.g., the effect of mileage on price. Second, the knowledge of similar examples, e.g., prices of other Ford Focus models might be relevant. Third, important landmark values, like the freezing point of water, might provide a starting point for answering the last question.

A dominant paradigm for research in quantitative estimation has been *anchoring and adjustment*. This work has been done in domains ranging from information rich, real-world estimation (Northcraft and Neale, 1987) to impoverished guessing (Tversky and Kahneman, 1974, Tenenbaum, in press). In this paper, we look at estimation in knowledge-rich domains and in naturalistic contexts. In such situations, experts routinely use analogical estimation: drawing upon similar examples from their experience while estimating. We begin with a brief survey of the relevant literature including a catalog of estimation processes. Next, we present our theory of analogical estimation. We then

present KNACK, a computational model of the analogical estimation process. Next we present the results of verbal protocols collected for realistic estimation tasks: a used car salesman estimating price of cars, and an apartment realtor estimating apartment rents. These data indicate that analogical estimation is frequently used during quantitative estimation. We then conclude with future work.

## 2 Background

Much of the literature relevant to quantitative estimation has focused on its numerical aspects. Even though numerical aspects are important, depending upon one's expertise, one may recruit varying degrees of semantic (non-numeric) knowledge while making estimates.

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 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) indicating that anchoring is not a purely numeric phenomenon, but has semantic underpinnings. Mussweiler and Strack's 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. The final numeric estimate is generated based on the easily accessible knowledge (Higgins, 1996), so their estimate is heavily influenced by anchor-consistent knowledge. This line of argument proposes a semantic priming based explanation of the anchoring and adjustment phenomena. Epley and Gilovich (2005) argue that the standard "experimenter-provided" anchors behave differently from "self-generated anchors," and the former

are not very informative about the actual process of adjustment.

Brown and Siegler (1993) explored quantitative estimation in a three-phase experimental paradigm. 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). 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.

## 2.1 A Catalog of Estimation Strategies

Several types of knowledge and reasoning processes are involved in quantitative estimation. Here we present a list of estimation strategies generated by analysis of expert solutions to estimation problems and used in a computer program that solves such questions (Paritosh and Forbus, 2005). This is based on an analysis of all problems (n=44) on Force and Pressure, Rotation and Mechanics, Heat, and Astronomy from Clifford Swartz’s “Back-of-the-Envelope Physics” (Swartz, 2003). There are two distinct processes:

**1) Direct Estimation:** This involves attempting to generate a quantitative estimate directly. One might know the value for the quantity sought, or one might have access to a salient landmark value that is known to be close to it. This is a common strategy when people are estimating dates, where they rely on temporal landmarks and dated period boundaries (Brown, 1990). People might also use *intuitive statistics* (Peterson and Beach, 1967) to extrapolate from known values.

**2) Transformations:** Failing direct estimation, the next step might be to transform the problem into other, possibly easier problems for which the answer might be known. These transformations can be divided into two types:

- Domain general transformations: For example, using a similar situation (analogy), using information about prototypes (ontology), decomposing an object using part-whole structure (mereology), using rates and averages (density).
- Domain specific transformations: Using physical laws, or rules of thumb that are specific to the domain.

An interesting aspect of estimation is its robustness with respect to amount of knowledge involved. Next, we present

our theory of how similarity and analogy are used in estimation.

## 3 Analogical Estimation

Analogy lets us make inferences from a better known example to a lesser 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. For example, when trying to estimate the rent for an apartment, one might retrieve from memory a similar apartment in the same neighborhood. The value from the analogical reminding serves as an *analogical anchor*. As a first pass, this analogical anchor is evaluated for its plausibility for the value sought. Analysis of the comparison between the problem and the reminding provides the grist for computing adjustments from the anchor to improve the estimate: for example, one might notice that the apartment that they were reminded of is smaller, and is in a slightly less desirable location. In this example, there are two causal assumptions about apartment rents:

- 1) A larger apartment has higher rent, all things being equal.
- 2) The more desirable the location, the higher is the rent, all things being equal.

These facts suggest that the estimate of rent should be more than the rent of the reminded apartment. Just how much more? The effect of location on rent can vary, and in some neighborhoods, it might be stronger than others. At this point, one can use other examples to determine just how strong that effect is. We call these adjustments based on causal knowledge *causal adjustments*. The final estimate is generated by adjusting the analogical anchor to reflect the causal adjustments. In this section, we explain in detail each of the steps above.

### 3.1 Analogical Anchors

Analogical estimation begins with searching and retrieving from memory other examples that are similar in ways to warrant being plausible estimates for the quantity sought. The reminders retrieved could be specific exemplars, or generalizations (Kuehne et al, 2000). The value of the quantity sought in the reminding is an analogical anchor. Analogical anchors are similar to *self-generated* anchors (Epley and Gilovich, 2004) in the sense that they are generated by the subject spontaneously as they solve the estimation problem. An example of a self-generated anchor is the freezing point of water while estimating freezing point of vodka.

However, there are two important differences between self-generated and analogical anchors: 1) the specific stimuli used in studies on self-generated anchors were designed to activate one strong anchor across subjects, and 2) self-generated anchors could be salient points on the dimension, irrespective of their relevance to the current problem. When one’s knowledge of the domain of estimation is sparse, they will recruit any salient points on the dimension to guide

their estimation. Most of the self-generated anchors fall in this category. However, with more experience in the domain, one might have access to a number of similar situations, possibly richly represented with causal knowledge and relationships between quantities. These are analogical anchors.

### 3.1.1 Quantities in Similarity

In order to use analogies to make numeric estimates, our analogical matching algorithms should be sensitive to quantities in the first place. 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. 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. One limitation of SME, 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 adequately. In most of these models, numbers are treated like symbols, so 99 and 100 are as similar/different as 99 and 10000. Models in case based reasoning (Ashley, 1990; Leake, 1996) that use numeric information employ ad hoc similarity metrics such as Euclidean distance that are not psychologically grounded. In order for a model of similarity to be useful in analogical estimation, we have following constraints on the retrieval and matching processes:

1) **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. (a large bird) should remind me of other large birds.

2) **Matching:** Similarity between two quantities should be computed and combined together in a cognitively plausible fashion, which amounts to answering: 1) How to compute similarity along a dimension? And 2) How to combine similarities along different dimensions?

A solution to these problems is symbolic encoding of the quantitative facts (Paritosh, 2003). Based on evidence from linguistics and psychology, we argued 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. CARVE (Paritosh, 2004) is a computational model for generating symbolic representations of quantity. Augmented with the symbolic representations generated by CARVE, we get better retrievals and matching of descriptions involving quantities.

### 3.1.2 Checking the plausibility of analogical anchor

The similarity between two objects doesn't necessarily warrant the inference that values of all the quantities for two objects are similar. For example, two similar basketball players might have similar height, but not necessarily two professors. This notion of what features can be inferred from a similar example was called *projectability* by Goodman (1955/1983). Projectability is based on centrality of the feature (Hadjichristidis et al, 2004). A feature is central to the extent that other features depend upon it. In the above example, height is central to basketball players, but not to professors. We have operationalized this notion of centrality as the structural support (Forbus et al, 1997) of the inference in computation of similarity using the SME.

### 3.2 Causal Adjustments

A key component of expertise is an understanding of the underlying causal structure of the domain. An important type of causal relationships is qualitative proportionalities (Forbus, 1984). Qualitative proportionalities indicate a monotonic relationship between two variables. These are useful for numeric estimation as they provide the ordinal direction for adjustment, e.g., a larger apartment has a higher rent, all else being equal. In verbal protocols presented in the section 5, we find that people commonly refer to such qualitative proportionalities while estimating. We call such adjustment based on qualitative proportionalities as causal adjustments.

However, it is not at all clear how to figure how much to adjust, as the qualitative proportionality only indicates a monotonic functional relationship between two variables, and does not tell us anything about the strength of this relationship. Let's suppose that the estimation problem involves two quantities:  $x$  and  $y$ , and that the unknown quantity we are trying to estimate is  $y$ . Further, we are given that there is a positive qualitative proportionality between  $x$  and  $y$ , i.e.,

$$y = qprop^+(x)$$

where  $qprop^+$  indicates a monotonically increasing function. Given this, if we were reminded of a similar situation, where the qualitative proportionality also holds true, and value of both quantities,  $x^*$  and  $y^*$  are known. Based on this, we can conclude if  $y$  will be more or less than  $y^*$ , as a result of the monotonic dependence.

$$sign(x-x^*) = sign(y-y^*)$$

At this point, we cannot conclude anything about how much more or less  $y$  is than  $y^*$  without making assumptions about the nature of the function  $qprop^+$ . However, if we know a few more examples where this qualitative proportionality is valid, i.e., data points on this function, we can use that to approximate the dependence by fitting a curve over those points. Let's assume we can recall a small set of situations where this proportionality is valid,  $\{(x_i, y_i)\}$ . Based on these, we can obtain an approximate estimate of the dependence between  $y$  and  $x$ ,

$$y = Q^*(x)$$

The suggested adjustment based on this approximation is,

$$adjustment = Q^*(x) - y^*$$

So, the causal adjustment is obtained by using an approximate estimate of the functional dependence between the quantities. The error in causal adjustment then is the discrepancy between this estimated dependence and that exists in the world. So, if one falsely believes that there is a strong relationship between two variables, then they are likely to produce a causal adjustment that is higher than needed. As opposed to the insufficiency results for adjustment, we expect errors in causal adjustments to be based on people's understanding of qualitative proportionalities in the world, and thus causal adjustments need not be insufficient. There is evidence to support that people can and do estimate correlations with very few samples, of the order of five (Kareev, 1997). We would expect causal adjustment to be affected by systematic biases in detecting such correlations.

### 3.3 Adjustment based upon non-alignable features

Comparison between the reminding and the problem might reveal features that are present in one but do not have a corresponding feature in the other (Markman and Gentner, 1997). For example, one might retrieve a similar apartment, but that one whose rent includes parking space. This is a sub-problem of the original estimation problem that is solved independently using the same mechanisms, e.g., one will invoke analogical estimation for the parking space.

## 4 KNACK: A Computational Model of Analogical Estimation

In this section, we present Knack, a computational model of analogical estimation. Figure 1 shows a high level description of Knack's algorithm. Knack's experience consists of a case library, a set of examples. An estimation problem is presented to Knack as a case, a set of predicate calculus expressions that represent all the information in the problem. Knack retrieves a few examples from the case library that are most similar to the problem at hand. The best reminding is used to generate the analogical anchor. The projectability of this inference is determined by looking at the structural support returned by SME. At this point, we extract all the aligned causal relationships that involve the quantity sought. A linear regression is performed for all the retrieved data points for each causal relationship. This gives us an approximate sense of the strength of the causal relationship. We compute adjustments for each causal relationship based on this approximate fit generated by linear regression. If the fit violates the expected qualitative relationship, then the adjustment suggested by this relationship is ignored. All valid causal adjustments are added to the analogical anchor to generate the estimate.

### 4.1 Estimating Basketball Statistics

To illustrate the above ideas, we report results from an experiment in the domain of estimating basketball player

1. Retrieve similar examples ( $n=5$ ) from memory
2. Select the most similar example's value as the anchor
3. Check if this is a plausible anchor by computing projectability
4. Find all causally connected quantities from the common causal structure in the retrieved examples
5. For each causally connected quantity
  - a. Compute adjustment via linear fit with the retrieved examples
  - b. Check adjustment with expected directionality of causal relationship
6. Apply all applicable adjustments to the anchor to generate the estimate

Figure 1. The KNACK algorithm

statistics (e.g., Points per game, Assists per game, height, etc.). This domain was chosen because there is a host of numeric information easily available, and there are interesting causal relationships between quantities, e.g., being tall helps to rebound. We selected thirty players such that they were reasonably different, six from each of the five positions on the court. We built a case library in which each basketball player was represented as a case. The average case had twelve facts, including four qualitative proportionalities, e.g., minutes per game is qualitatively proportional to points per game.

#### Method

We compared Knack to baseline analogy by running two trials. The baseline trial makes estimates by choosing the value for the dimension on the player selected by MAC/FAC as the best reminding. The Knack trial utilized CARVE to enrich the cases with symbolic representations for the quantitative facts. On an average, this added ten facts to every case. For example, CARVE generates the following qualitative representation for each quantitative fact:

Quantitative fact	Qualitative representation
(seasonThreePointsPercent JasonKidd 0.404)	(isa JasonKidd (HighValueContextualizedFn seasonThreePointsPercent BasketballPlayers))

In both trials, the facts mentioning the sought after dimension were filtered out of the question case. The trials were conducted in round robin format in which estimates were recorded for every player and every dimension.

#### Results

We present the comparison of error in estimates generated using baseline analogy and Knack. Knack's estimate are significantly more accurate ( $p < 0.05$ ) for four out of six dimensions across all players. Although the error for assists per game appears to be higher for Knack, the difference is not significant. Similarly, there is no significant difference in errors for free throw percentage. The free throw percentage dimension was not causally related to any other quantities, and the assists per game is highly variable across our dataset. This is because our representation implies that these dimensions are not causally central.

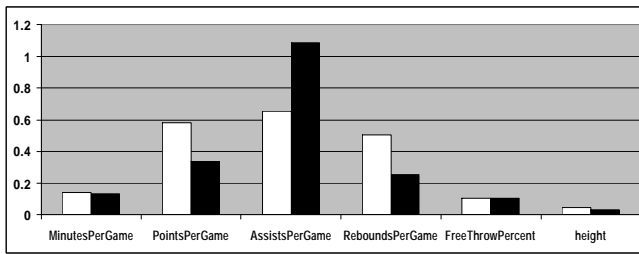


Figure 2: Comparison between normalized mean error,  $ABS(estimate - value)/value$  of estimates by dimension. White is baseline error and black is Knack's error.




Figure 3: Comparing under adjustments, over adjustments, and wrong adjustments made by Knack over all the estimation problems.

When looking at the amount and direction of the adjustment, we consider no adjustment to be an under adjustment. Qualitative proportionalities are directional, therefore Knack only made causal adjustments for points per game, rebounds per game, and assists per game. Knack handles the contradiction between a computed adjustment direction and the sign of the qualitative proportionality by ignoring it. This leads to a systematic under adjustment for these dimensions.

## 4.2 Discussion

Knack demonstrates how similar examples can be used to find analogical anchors in quantitative estimation tasks. These analogical anchors are similar to the self generated anchors studied by Epley and Gilovich (2005). They found that with forewarning and incentives subjects could overcome the insufficient adjustment bias. The Knack model hypothesizes that under adjustment is more likely when the subject is less confident in the nature of the adjustment. One way in which Knack could model the increased effort in overcoming the bias would be when faced with contradictory adjustment directions, to retrieve more and more examples until the qualitative proportionality was satisfied. Knack is consistent with Mussweiler and Strack's (2001) claim that anchors cause subjects to activate anchor consistent knowledge. The symbolic encoding of the anchor as a plausible answer will bias the retrieval of examples that are consistent with it.

## 5 Verbal Protocols of Expert Estimation

To observe how experts utilize similarity and causal relationships in real world estimation tasks, we conducted a protocol analysis of experts doing realistic estimation tasks. The goal of this study is to determine the extent to which experts use analogical estimation, specifically if they use analogical anchors and make causal adjustments while estimating. We interviewed two experts in two different domains, an employee at a used car dealership with two years of experience and a apartment realtor with five years of experience. We constructed numerical estimation questions by taking items off of public listings within areas of their expertise, online advertisements for housing rentals and used cars, and removing the asking price.

### 5.1 Method

After performing some warm up exercises as recommended by Ericsson and Simon (1993) to increase the participants willingness to reason aloud, we would present each problem. The subjects were given as much time as they wanted to answer each question, and they were occasionally prompted with questions such as "What are you thinking about right now?", if they remained quiet too long. Each question was a complete listing of a car or an apartment with the price removed. Apartment listings were from the website craigslist.org containing all details such as size, location, utilities, etc. This is the same information an apartment seeker would have access to in their preliminary search. Car listings were taken from a popular used car website carmax.com. Each description contained a picture of the specific car and a standard format containing relevant information about the car. After each estimate the participants were asked to explain their answer. The apartment trial consisted of eight questions, while the used car trial consisted of seven questions.

### 5.2 Coding Scheme

We coded the protocols for three aspects of analogical estimation:

#### 1. Analogical reminders

Explicit references to remembered prototypes of a class, or specific instances that were similar to the problem, e.g., "This [Lakeview apartment] would go for \$700-750 in Rogers Park" and "These [cars] are just shy of \$30,000 brand new."

#### 2. Causal adjustments

Explicit references to other causal quantities during the estimation process, e.g., "You know [parking spaces] are worth more in Lakeview" and "These [cars] are particularly hot right now because of higher gas prices."

#### 3. Non-alignable features

Explicit adjustments based on features present in one of the cases, e.g., "If [the Cadillac Escalade] is black it is 1,000 dollars more" and "I'm going to raise [the estimate] a little, I was not thinking about the deck."

Table 1: Number of analogical estimation occurrences apparent in each trial.

Analogical Estimation Aspect	Cars (n=7)	Apartments (n=8)
Analogical reminders	7	11
Causal adjustments	11	7
Non-alignable adjustments	5	12

The results indicate that analogical estimation is a common strategy used by experts solving estimation tasks.

## 6 Conclusions and Future Work

Analogical estimation is a key part of real-world quantitative estimation. We presented a theory and a computational model of analogical estimation. Analogical estimation involves using analogical anchors, and making causal adjustments based on inferred strength of causal relationships. Consequently, we expect that anchoring and adjustment will show a strong effect of experience with other examples in the domain: 1) the most similar examples will provide the anchors, and 2) adjustment need not be insufficient, but will mirror the strength of the causal relationships that follow from the subject's experience. More psychological experiments have to be done to verify these predictions.

## 7 Acknowledgments

The authors would like to thank Ken Forbus for insightful comments on the drafts. This research is supported by the Computer Science Division of the Office of Naval Research and DARPA.

## 8 References

- 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. (1990). Organization of public events in long-term memory. *Journal of Experimental Psychology: General*, 119, 297-314.
- 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*. 100(3). 511-534.
- Chapman, G.B. and Johnson, E.J. (1999), Anchoring, activation and the construction of values. *Organizational Behavior and Human Decision Processes*, 79(2), 115-153.
- Ericsson, K. A., Simon, H. A. (1993). Protocol Analysis: Verbal Reports as Data. MIT Press, Cambridge, MA
- 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.
- Forbus, K., Gentner, D., Everett, J. and Wu, M. (1997). Towards a computational model of evaluating and using analogical inferences, In *Proceedings of Cognitive Science Conference*.
- 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.
- Griffiths, T. L. and Tenenbaum, J. B. (in press). Optimal predictions in everyday cognition. *Psychological Science*.
- Hadjichristidis, C. Sloman, S., Rosemary, S. and Over, D. (2004). Feature centrality and property induction. *Cognitive Science*, 28 (2004), 45-74.
- 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.
- 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.
- Kahneman, D. (1992). Reference points, anchors, norms, and mixed feelings. *Organizational behavior and Human Decision Processes*, 51, 296-312.
- Kareev, Y., Lieberman, I., and Lev, M. (1997). Through a Narrow Window: Sample Size and Perception of Correlation, *Journal of Experimental Psychology: General*, 126, 3, 278-287.
- Kuehne, S., Forbus, K., Gentner, D. and Quinn, B.(2000) SEQL: Category learning as progressive abstraction using structure mapping. In *Proceedings of Cognitive Science Conference*.
- Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory & Cognition*, 24(2), 235-249.
- Mussweiler, T. and Strack, F. (2001). The semantics of anchoring. *Organizational Behavior and Human Decision Processes*, 86(2), 234-255.
- 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. (2003). A Sketch of a Theory of Quantity, In *Proceedings of the 17th International Workshop on Qualitative Reasoning*, Brasilia, Brazil.
- Paritosh, P.K. (2004). Symbolizing Quantity. In *Proceedings of the 26th Cognitive Science Conference*, Chicago.
- Paritosh, P.K. and Forbus, K.D., (2005). Analysis of Strategic Knowledge in Back of the Envelope Reasoning, In *Proceedings of the 20th National Conference on Artificial Intelligence*.
- Peterson, C.R., and Beach, L.R. (1967). Man as an intuitive statistician, *Psychological Bulletin*, 68(1), pp 29-46.
- Tversky, A., and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases, *Science*, 185, pp 1124-1131.