

Praveen K. Paritosh

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

Abstract

Quantities are ubiquitous and an important part of our understanding about the world - we talk of engine horsepower, size, mileage, price of cars; GDP, population, area of countries; wingspan, weight, surface area of birds, and so on. In this paper, we present a sketch of a theory of quantity - cognitively sound representations and principles for generating those representations. We present evidence from psychology, natural language, and ecological constraints to argue for a cognitively plausible representation of quantities. We then propose a general principle of how to make the necessary and relevant distinctions. Structured models of retrieval, similarity, and generalization, and in general models involving symbolic representations, do not handle quantities adequately. That is an artifact of poor representations of quantity, and we believe that the representations proposed here will make these models more quantity-aware. This investigation is at the intersection of qualitative reasoning, cognitive psychology, and linguistics, and builds on existing evidence in these fields to potentially contribute to the understanding of quantities in all the three.

1 Introduction

The notion of quantity is quite broad, and there is a substantial literature in psychology, linguistics and qualitative reasoning (QR) on many different aspects of it. The psychology of perceptual quantities, those for which we have direct sensory measure, (the psychological literature refers to them as *dimensions*) like brightness, loudness, etc., has been studied in detail¹, but conceptual quantities like price of computers, GPA of students, etc., have not been studied as much. Most of the research from the decisionmaking community and the case based reasoning community [Leake, 1996] that does study conceptual quantities employs either metric space or feature vector models. On the other hand, the structured models of similarity and generalization, which have strong converging psychological evidence, do not handle quantities adequately. In linguistics, there is relevant work on the nature of dimensional adjectives like large, small, hot, etc. The QR literature contains many different proposals for qualitative representations of quantity [Forbus, 2003]. Each of these efforts have different goals. In this paper, we take a closer look at what our representations of quantities contain, guided by cognitive and linguistic evidence, and ecological

constraints on our knowledge of quantities. More specifically, we address the following two fundamental questions:

- 1. What do our (cognitive) representations of quantities look like? Or, what representational machinery is needed for quantities?
- 2. How are these representations built with experience?

These are questions about how cognition works, as well as about how the world is organized. This is opposed to Bierwisch (1967), who argued that dimensional adjectives (like hot and cold) do not represent "properties of surrounding world in the broadest sense, but rather certain deep seated properties of the human organism and the perceptual apparatus." This distinction is important, as many of representations of quantity (for example, in scientific domains, and thus those in QR) tend to have the underlying perspective of representing properties of the world in some optimal fashion. The Boiling Point, for example, seems to be more a property of the world than the notion of Expensive, which seems more variable and harder to formalize. For example, consider 'temperature' - we find adjectives like cold, tepid, lukewarm, warm, hot, which are quite different from the necessary and relevant [Forbus, 1984] distinctions. Rather than being redundant and/or informationally sub-optimal, we argue that those linguistic distinctions play an important role in our representations of quantity and their development.

Although related, the answers to the two questions above are fairly independent. To address the former, we draw arguments from existing evidence in cognitive science, natural language, and the constraints on our representation arising from the reasoning tasks and the nature of variation of quantities in real world. In section 3, we argue that quantity spaces are part of solution, but more is needed. As tasks broaden, e.g., doing similarity based reasoning, the sources of constraints on the qualitative representations are different. The quantity space representation [Forbus, 1984] has the expressive power that our representations of quantity seem to have, and we extend and provide cognitive and linguistic evidence in its support. The story, though, is not complete, until we have a principled way of finding the important symbols on the space of values that the quantity can take. Section 4 tackles the second question above, and proposes a general principle underlying the necessary and relevant distinctions. Section 5 presents a plan for implementing and testing the ideas presented here. The last section

¹ Many interesting insights can be found in this work – integrable/separable distinction, analytic/holistic perception, developmental trends – see Tighe and Shepp, 1983 for a collection of articles on these issues.

concludes with a few other relevant questions that have not been addressed here.

2 Background and Motivation

This section presents relevant background in qualitative representations and structured models of retrieval, similarity and generalization – to both of which this theory will make potentially useful contributions – providing cognitively sound qualitative representations, and extending the structured models to include effects of quantities, something currently ignored in those models.

2.1 Qualitative Representations of Quantity

One of the goals of qualitative reasoning research has to understand human-like commonsense been reasoning without resorting to the preciseness of models that consist of differential algebraic equations and parameters that are real-valued numbers. There is a substantial body of research in QR that has shown that one can, indeed, do a lot of powerful reasoning with less detailed and partial knowledge. Qualitative reasoning has explored many different representations: status algebras (normal/abnormal); sign algebra (-, 0, +), which is the weakest representation that supports reasoning about continuity; quantity spaces, where we represent a quantity value by ordinal relationships with specially chosen points in the space; intervals and their fuzzy versions; order of magnitude representations; finite algebras, among others. The representations differ in the kind of distinctions that they allow us to make. To echo the questions raised in the introduction, we are interested in finding a cognitively sound representational framework for these distinctions, and principles for finding the distinctions that we do and should make.

Our answer to the first question raised in the introduction is that the quantity space representation, augmented with distributional information, accounts for observations and existing evidence from psychology and linguistics. The current evidence does not conclusively prove or disprove this claim, and we feel that bridging this gap between QR and cognitive science will be a contribution to both fields. Our answer to the second question is the first attempt to come up with a general theory of what distinctions to make. Sachenbacher and Struss (2000) attacked a similar problem – they were interested in finding the right distinctions given the reasoning task. Here we are more concerned with cognitively plausible distinctions - for example, the distinctions that are made in natural language on the space of values that the quantity takes.

2.2 Retrieval, Similarity and Generalization

There is converging psychological evidence for structured models of retrieval, similarity and generalization. In contrast, feature vector models [Leake, 1996] employ ad hoc similarity metrics. Below we give a very brief introduction to the idea of structure mapping, and then present the shortcomings of structured models in regards to quantities.

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). Each mapping is a set of correspondences that align particular items in the base with items in the target, and candidate inferences which are statements from the base that are hypothesized to hold in the target by the virtue of these correspondences. 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 data structure, the content vector, which is a summary of relative frequency counts of various symbols in the structured representation. The dot product of content vectors of two structured representation provide a rough estimate of their structural match. SEQL [Skorstad et al, 1988; Kuehne et al, 2000] provides a framework for making generalizations based on computing progressive structural overlaps of multiple exemplars.

Representing them as numbers does not go far in being useful – for example, our models of retrieval (MAC/FAC), similarity (SME) and generalization (SEQL) do not care much about quantities represented such. The way quantities are implicated in these processes –

Retrieval: Just as Red the symbol 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. Then it will show up in content vectors, and contribute to the dot product.

Similarity: A model of similarity that is sensitive to quantities will explain how quantity values can make two structured descriptions that have similar amount of structural overlap more or less similar (for example, in SME, 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). Which means, answering -1) how to compute similarity along a single dimension, and 2) how to combine the similarity along different dimensions in computing overall similarity of two cases. Also, what inferences about quantity values should be sanctioned by structure-mapping? or example, inferring how much I will spend on a conference trip based on a previous conference trip.

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, their shelters, etc.

All of the above are currently not supported in SME, MAC/FAC and SEQL. A large part of this deficiency,

we feel, is due to poor representations of quantity. A symbolic and relational representation of the kind we propose here would automatically make these models more quantity-aware.

3 Cognitive Representational Machinery

Our knowledge about quantities is of various kinds – we talk of Expensive and Cheap things, we know that basketball players are usually Tall, we know that Boiling point of water is 100 degree Celsius. In this section, we present and argue for a cognitively plausible representation of quantity. There are three subsections – 3.1 organizes arguments for what must be contained in our representations of quantity around various constraints, 3.2 presents the proposed representation, and 3.3 discusses some implications of this representation.

3.1 Constraints

Representations do not arise in vacuum – they are molded by the kinds of reasoning tasks we perform with them (reasoning constraints), the underlying reality of the things we are trying to represent (ecological constraints), and how we perceive this reality (psychological constraints). Based on these, and scattered pieces of evidence from psychology and linguistics, we argue that our representational machinery for quantities must contain partially (or possibly totally) ordered symbolic reference points (*a la* quantity space), and distributional information about the quantity (or an informational equivalent thereof).

3.1.1 Reasoning Constraints

The three distinct kinds of reasoning tasks involving quantities are –

1. Comparison: These involve comparing two values on an underlying scale of quantity (or dimension²), e.g., "Is John taller than Chris?" 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 and judging the larger of two large things than the smaller of two large things - e.g., subjects are faster and more accurate at interpreting "A whale is larger than an elephant" than "An elephant is smaller than a whale." 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.

2. Classification: These involve making judgments about whether a quantity value is equal to, less than or greater than a specific value³, e.g., Is the water boiling?, Will this couch fit in the freight elevator?, Is he below the poverty line?, etc. Usually, such classifications involve comparisons with interesting points on the space of values that a quantity can take, moving across which has consequences on other, different aspects of the object in concern. The metaphor of *phase transitions* describes many of such interesting points, although such transitions in everyday domains are not as sharply and well defined as in scientific domains (consider Poverty line versus Freezing point). We talk about this more later in the next section.

3. Estimation: These involve inferring a quantitative/numeric value for a particular quantity, e.g., How tall is he? What is the mileage of your car? This is the activity that has the strongest connection to quantitative scales – one can go a long way in accounting for the above two without resorting to numbers, but estimation involves mapping back to numbers [Subrahmanyam and Gelman, 1998]. Knowledge of interesting points on the scale might play an important role in estimation, for example in providing *anchors* to *adjust* from [Tversky and Kahmenan, 1974].

These tasks are not completely distinct from each other – classification involves comparison, and estimation might be used in the service of classification. Two interesting aspects of our representations follow from these constraints –

1. Our representations must keep track of interesting points on the scale of quantity, to classify, as well as to estimate.

2. Labels like large ease making comparisons, as they setup implicit ordinal relationships (it is larger than the expected norm), which seem to be references to the underlying distribution of the quantity values.

3.1.2 Ecological Constraints

Quantities vary in a different fashion than, say, nominal attributes. Our representational framework must be capable of capturing the interesting ways in which a quantity varies in real-world instances of it. Below we present two different kinds of constraints on values a quantity can take –

1. 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. References to Tall and Short men, then seems to be a reference to an underlying distribution of heights of people. A popular account of dimensional adjectives (e.g., "Flamingo is a large bird") is that it establishes a comparison to an underlying categorical norm [Rips, 1980; but see Kennedy, 2003]. But it seems more than just reference to the norm; anything greater than the norm

²Consider "The space telescope is longer than it is wide." These cross-dimensional comparisons can get quite complicated to interpret, e.g., "The Sears tower is as tall as the San Francisco Bay Bridge is long" does not literally mean that Height(Sears Tower) <= Length(San Francisco Bay Bridge). See Kennedy (2001) for an analysis and implications of such comparisons.

³ Or corresponding judgments involving intervals.

is not large or high - it also implicitly takes into account the spread of the distributions. More than just the norm, 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. We describe the most compelling of such studies here - refer to Peterson and Beach, 1967; Fried and Holyoak, 1984; Kraus et al, 1993; Ariely, 2001, among others, for more. Malmi and Samson (1983) presented subjects with one hundred three-digit numbers, which they were told were SAT scores of two different groups (named PIM and DAP). Each "SAT score" was displayed as either of PIM or DAP, and the three-digit number. Even when the numbers were displayed for merely 0.5 seconds, subjects accurately estimated (within 95% confidence interval of the stimulus mean) the mean for both PIM and DAP samples in the case of both normal, skewed distributions and bimodal distributions. The last one suggested that subjects might be storing more than just a running mean, and so they tested the subjects for how accurately could they reproduce the entire frequency distribution of the sample, and the subjects were able to reproduce the distributions qualitatively, as well as quantitatively.

Surprisingly, the next question of how we partition these distributions has not been raised at all. 2. Structural Constraints: A quantity is also constrained by what values other quantities in the system take, its relationship with those other quantities, the causal theories of the domain; in general, the underlying structure of representation⁴. 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 are very interesting, as they represent the underlying mechanism, or the causal story of the object. As we move along the space of values a quantity can take, it is possible that we transition into a region where the underlying causal story is different (e.g., ice starting to melt, at the freezing point), which induces extremely important and interesting distinctions of quality on the space of quantity. Much of the representations in QR involve such transition points.

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.

3.1.3 Psychological Constraints

The way in which we acquire knowledge about quantities leads to making the following distinction between two classes of quantities –

1. Perceptual: We have a direct sensory measure of these attributes, e.g., brightness, loudness, size of things that we see, etc.

2. Conceptual, or Abstract: No direct sensory perception of these attributes, usually expressed using numbers, e.g., prices of things, mileage of cars, gross domestic products of countries, the size of countries, the GPA of a student, the clock speed of a computer.

Another distinction between perceptual and conceptual quantities is that the former has been studied in much more detail. Sometimes, though, it might not be clear whether a quantity's underlying model in the person's head is perceptual or conceptual, as one might *perceptualize*⁵ even the most abstract quantities (and vice versa) – e.g. the size of a country as the size on the map. This paper focuses on the conceptual dimensions, where rich structured representations and higher-level cognitive processes might play a bigger part than that in perceptual dimensions.

3.2 Proposed Representation

Based on the observations in section 3.1, here we propose that our representations must contain symbolic reference points, and distributional information.

3.2.1 Symbolic references to quantity

A partially, or possibly totally ordered set of symbolic reference points forms the *quantity space* [Forbus 1984]. Any value on the scale can then be represented via ordinal relationships to these symbolic reference points. Quantity space is the minimal representation that supports variable resolution. The symbolic and relational nature of this representation automatically makes it much more useful in our (structured/symbolic) representational framework. In the original formulation of quantity space, these symbols are limit points, those points where different processes/model fragments become active or de-active. We will relax that constraint, in the discussion to follow, see what other kind of reference points are needed.

The psychological reality of such special 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 stimuli is judged closer to a reference stimuli (e.g., the color offred to basic-red; the number 996 to 1000) than otherwise, while such relationship between two nonreference stimuli is symmetric. Existence of landmarks to organize spatial knowledge of the environment [Lynch, 1960], similar asymmetries [Holyoak and Mah, 1984 among others], and effects on encoding and retrieval [Ferguson and Hegarty, 1994] has been

 $^{^4}$ 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.

⁵ See, for example, Gattis, 2002, for a set of articles on spatial metaphors/representations of abstract thought.

shown. Other relevant psychological studies that support the existence of reference points come from categorical perception [Harnad, 1987], magnet effect in speech perception [Kuhl, 1991] and sensitivity to landmarks [Cech and Shoben, 1985]. Brown and Siegler (1993) proposed the *metrics and mappings* framework for real-world quantitative estimation. They make a distinction between the quantitative, or metric knowledge (which includes distributional properties of parameters), and ordinal information (mapping knowledge). The two main types of symbols in our quantity spaces are –

1. Structural Limit Points: 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. Distributional Partitions⁶: Symbols like Large and Small, which arise from distributional information about how that quantity varies.

These two cover a large part of symbolic references to quantity in language. We have not found conclusive evidence in existing literature that argues for symbolic reference points/intervals in our representations, or talks about the kind of symbols. One of the goals of this project is to set this representation on firmer grounds, bridging evidence from psychology and linguistics.

3.2.2 Distributional information

A distribution is a summary of how the quantity value varies. This is the information that we accumulate based on our interaction with concrete exemplars, and specific quantity values. These distributions provide the grist for the symbolic distributional partitions. If we take the existing psychological evidence to mean that we accumulate distributions of quantity values, the unanswered questions are –

- 1. What is the class of objects for which we compute and store distributions? So, do we have a distribution of all lengths, or length of vehicles, or length of cars, or length of sedans? This is similar to the question of how to determine the comparison class for an adjective [Kennedy, 2003]. Clearly, some of such distributions are computed dynamically based on the context, [Rips and Turnbull, 1980; Staab and Hahn, 1998] – it yet remains to be established the ones that are stored along with our representations.
- 2. How do the symbols on the space of quantity (both structural and distributional partitions) map on to these distributions?

3.3 Implications

Informal analysis of symbolic references to quantity in natural language provides support for the representation proposed above, but there are interesting differences. There are some interesting issues lurking here –

Intervals versus points: Why is it that physical and scientific domains (and thus qualitative reasoning)

finds the transitions, the points (freezing point, boiling point), more interesting; as opposed to language, where it seems that most of the references are to intervals (cold, hot), even when the transitions are sharper (e.g., wet/dry)? A plausible conjecture is that intervals let us talk about the quantity without making commitments to where exactly the transitions happen.

Crisp versus soft transitions: In most scientific domains, most of the transitions (e.g., freezing point) are crisp, but many of such transitions in other everyday domains seem to be *softer*. We suspect that there might be a psychological distinction between the crisp and soft transitions – discrimination should worsen across the soft transitions (tall/short) as opposed to crisp transitions (wet/dry).

4 Necessary, relevant, and more distinctions

The last section argued for the quantity space representation – that the symbolic reference points provides a way to make the distinctions on the space of values of a quantity that we make. Here we address the second question raised at the beginning of this paper – which amounts to asking – how many symbols do we use, and where and how do they map on the space of quantity values⁷?

4.1 Structural Limit Points

Structural limit points are a generalization of the idea of limit points introduced in QP theory. One should only make the *necessary* and *relevant* qualitative distinctions, QP theory advises us. In the domain of processes, QP theory provides the intuition for these distinctions: *where things change*, i.e., different processes and/or model fragments get de/activated, e.g., Freezing Point and Boiling Point of a liquid. Is there a general principle that provides these distinctions for more than just dynamical processes?

One can always partition the quantity space arbitrarily – so, one could have an ad hoc rule that said that we'll always divide the space between the minimum and maximum into three parts – high, medium and low⁸. That would mean that we are suggesting that there are some partitions that are more *natural* than others. Some features of the natural partitions –

Right level of granularity: Freezing Point and Boiling Point might be fine for reasoning about physical behavior, but if one is talking about shower water, then more distinctions like Cold, Body Temperature, Warm and Scalding Hot might be more appropriate.

Structurally predictive: of other properties of the system, e.g., Poverty Line, Lower Class, Middle Class, Upper Class.

⁶ Interesting asymmetry here – Most of the distributional information is symbolized as intervals, and not points.

⁷ The question being asked here is what and where are the distinctions to be made, not what to label those distinctions.

⁸ For example, the Fuzzy logic community does something in the same spirit.

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 structure⁹ therein. Processes (as in QP theory), are a special case of these structural bundles (where the key relationships are of causality and influence) for the class of dynamical physical systems. Thus, the key idea is –

The necessary and relevant qualitative distinctions correspond to discontinuities in the underlying reality as captured by the structure in the representation.

Let us look at an example – consider people's income. Poverty line, lower class, middle class and upper class define changes of quality on the space of income, as we expect that many other aspects of people – their lifestyle, the amount of time/money they spend on entertainment, education, the kind of vacations they have (or do not), the family and social climates in which they live, their expectations and relationships to the rest of the social structure, etc. changes as we move across these interesting partitions of the scale of income.

Consider the size of dictionaries (as measured in number of pages, volume, or weight). There seem to be at least three meaningful distinctions of quality that might be projected on to size – pocket, table-top, and library-sized dictionaries. The underlying reasons/stories for these three types of dictionaries are quite different – the key aspect of the pocket dictionary is portability, and thus it has finer print, thinner pages, less detailed meaning, probably not much etymology and usage information, etc; while the key aspect of the library sized one is comprehensiveness, and thus it follows that it is larger, heavier, has a much higher number of entries and even arcane and obsolete words, etymologies, usage information, is well bound as it is big and thick, pages are tougher to stand more usage, etc. The table-top dictionary falls somewhere in between. On the dimension of size, thus, the distinctions of pocket, table-top and library-size, define interesting distinctions which have deep relationships to the underlying causal story, the underlying quality of dictionaries.

These changes of quality in the above two examples are reminiscent of phase transitions in physics/thermodynamics – just as a lot of underlying properties and the relationships that tie them together change as we move across the phase transitions. There are two types of phase transitions – first-order (sharp discontinuity, e.g., solid—liquid change), and secondorder (where one can continuously move from one phase to another, e.g., magnetization) [See Sethna, 1992 for an introduction, and and Gunton et al, 1983 for more detailed explanation]. The structure of relationships is the analogue for equations of state that hold in a particular phase, and the crisp/soft distinctions are analogous to first-/second-order transitions.

4.2 Distributional Partitions

The importance of the structural limit points presented in the previous section is apparent – it is predictive of structural properties of the system, and thus quite useful in doing qualitative reasoning. Surprisingly, though, the language contains many references to quantity which look very different from the structural limit points or the intervals they might induce consider Large, Tall, Short, Expensive, etc. When we say a large flamingo, that seems to be reference to the distribution of sizes of flamingos and the fact that the particular flamingo we are looking at is larger than the norm. Such distinctions like small, medium, large, seem to be making cuts based on the distribution of values that the quantity takes, and the four most common distributions - uniform, normal, skewed, and Zipf, have different intuitions. An intuitive understanding of the normal distribution might be that there are fewer short and tall people than there are people of regular height (and also that the range of tall and short is larger than the regular size). The power law, or the Zipf distribution is extremely interesting, as it is extremely common – a meaningful norm for such distributions can not be defined. Are there some systematic ways in which people make cuts on a distribution they have abstracted? There is little known about this.

4.3 Implications

Most symbolic references to quantity have both a structural and distributional interpretation of them – so being Tall has structural consequences, for example, for a basketball player (or a gymnast). An interesting question is to see the interactions between these two types of partitions. When do we choose to use structural partitions, and when distributional? The answer has to do with the nature of the quantity. Some quantities are more causally central – i.e. more deeply affect other aspects of the system, than others (compare horsepower of a car to size of the door handles). In the class of examples that we are looking at, there will be a tendency to describe a quantity purely using distributional information if –

- 1. The parameter doesn't have deep causal connection to the rest of the system, or is not causally central (in terms of structured representation, has low *systematicity*) – height of poets as compared to height of basketball players.
- 2. There is not much of variation in the underlying structure (as far as is known in our representation) at all, e.g. size of adult male penguins.

⁹ The structure of relationships is an even more general notion than causality, spatial arrangement, connectivity.

We asked in the introduction about why language is full of dimensional adjectives like hot, cold, etc., which usually are distributional partitions; as opposed to structural limit points, which are the kind we find in scientific domains, and thus QR. Our conjecture is the three parts of our representations of quantity distributions, distributional partitions, and structural limit points, in that order, support each other. When one begins to learn a domain, the distributions are accumulated until we know enough to give them symbolic labels - and that helps us build the causal structures that then lead to the structural limit points. Dimensional adjectives also allow for the flexibility in usage and interpretations, making them their useful. Context sensitivity of our linguistically partitions of quantity might be captured by the representation of the class of examples in such a manner that we have all the context relevant aspects well represented in it.

5 Future Plans

We are just starting to implement the ideas presented above to support, test and refine them. SEQL will provide the basic framework for finding structural clusters. Following are the next steps in this research –

Building corpus of examples: We will build at least three different case libraries of examples rich in quantities and structured representations to test these ideas. The CYC knowledge base has information about countries taken from the CIA World Factbook. There are about thirty quantities whose numeric values are known for most countries (population, area, birth rate, etc.). We are adding structural knowledge about the relationships between these parameters.

Building the next generation SEQL

- 1. Bootstrapping: To find the interesting structural limit points, we cannot ignore the numbers altogether, and thus we need to a do a rough, first-cut symbolic partitioning for the quantities. In the previous section, we mentioned distributional and structural constraints. This is where the distributional constraints play a role. Purely by looking at the distribution of a single quantity, one can divide it into ranges (e.g., low, medium, high population).
- 2. Extending category representations and SEQL to assimilate and include distributional information.
- 3. Develop into an anytime/incremental version that incorporates all the ideas here.

Evaluation: Once we have completed the above, the goal is to see if we can use the above ideas to generate symbolic representations for the quantities. For testing the 'goodness/natural-ness' of these representations, we can -

- **1.**Compare the representations thus generated to experts' qualitative representations of the same quantities.
- **2.** Incorporate these representations into the cases, and see if we get better retrieval and similarity measurements as compared to human subjects.

3. Build an analogical quantity estimator, which will estimate a value for a quantity by retrieving a strongly similar example for which it knows the value. This will be a part of the Back of the Envelope Reasoner [Paritosh and Forbus, 2003].

6 Conclusions

The discussion above is an attempt to ground qualitative representations of quantity in cognitive and linguistic evidence. The distributional and structural limit points provide a symbolic and structural representation of quantity that has the potential to be quite useful in our framework, making both steps of MAC/FAC more sensitive to quantities. It will provide a principled way to partition the space of a quantity, telling us which quantities are more predictive (those that are more systematic, or structurally central) than others of structural properties. Relational representations of quantity (ordinal relationships with the structural limit points/ distributional partitions) automatically provide for a way for SME to take quantity into account while computing similarity, and also to make quantitative inferences sanctioned by analogical comparison. Furthermore, en route to generating these representations, we would extend SEQL to account for learning the sense of quantity, an important part of back-of-the-envelope reasoning.

Some very interesting questions that are raised that we did not cover in this paper – $\ensuremath{\mathsf{-}}$

- 1. How do the quantities/dimensions come to be in the first place? Presumably, part of learning a domain is *recognizing* the quantities in the first place (e.g., Clark, 1973, presents evidence that 'big' is easier than length and height, which are easier than width "semantic feature hypothesis").
- 2. Where exactly are the symbols that make up the quantity space and the ordinal relations between them stored in the generalizations, or each of the exemplars?
- 3. What are the dynamic aspects and what are the stored/static versions of these representations?
- 4. There is a class of landmarks which are special mainly because of their perceptual salience or memorability, e.g., birthdays [Shum, 1998], powers of ten [Rosch, 1975], or many of the spatial landmarks that we use [Lynch, 1960]. Are there any interactions between such landmarks and the structural/distributional partitions proposed in this paper?

Clearly, a full understanding of quantities is a major enterprise, but we are confident that the ideas here will form a part of it.

7 Acknowledgments

This research is supported by the Computer Science Division of the Office of Naval Research. The author would like to thank Ken Forbus, Dedre Gentner, Chris Kennedy, Lance Rips, Sven Kuehne, Karen Carney, Emmett Tomai, Sam Day and Rachel Goldsborough for insightful comments and discussion on the work presented here.

8 References

- Ariely, D. (2001). Seeing Sets: Representation by statistical properties. Psychological Science, 12(2), 157-162.
- Bierwish, M. (1967). Some Semantic Universals of German Adjectivals. *Foundations of Language*, **3**, 1-36.
- Banks W. P., and Flora J. (1977). Semantic and Perceptual Processes in Symbolic Comparisons. *Journal of Experimental Psychology: Human Perception and Performance*, 3, 278-290.
- Brown, N. R., & Siegler, R. S. (1993). Metrics and mappings: A framework for understanding real-world quantitative estimation. *Psychological Review*, **100**(3), 511-534.
- 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.
- Clark, E.V. (1973). Whats in a word? On the child's acquisition of semantics in his first language, In T. E. Moore (Ed), *Cognitive development and acquisition of language*, Academic Press, NY.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41, 1-63.
- Ferguson, E. L. and Hegarty, M. (1994). Properties of cognitive maps constructed from texts. *Memory and Cognition.* 22(4), 455-473.
- Forbus, K. D. (1984). Qualitative process theory. Artificial Intelligence, 24, 85-168.
- Forbus, K. D. (2003), Qualitative Reasoning, CRC Handbook of Computer Science and Engineering.
- Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19(2), 141-205.
- Fried, L. S., and Holoyak, K. J. (1984). Induction of Category Distributions: A Framework for Classification Learning. *Journal of Experimental Psychology: Learning, Memory,* and Cognition, 10, 234-257.
- Gattis, M. (Ed.) (2002) Spatial Schemas and Abstract Thought. Cambridge, MA: MIT Press.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gunton, J. D. San Miguel, M. and Sahni, P.S. (1983). The dynamics of first order phase transitions, In *Phase Transitions and Critical Phenomenon* (Eds. C. Domb and J. L. Lebowitz), Vol. 8, Academic Press, London.
- 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.
- Joram, E., Subrahmanyam, K. and Gelman, R. (1998). Measurement Estimation: Learning to Map the Route from Number to Quantity and Back. *Review of Educational Research*, 68, 413-449.
- Kennedy, C. (2003). Towards a Grammar of Vagueness. To be presented at the *Princeton Semantics Workshop*, May 17, 2003.
- Kraus, S., Ryan, C. S., Judd, C. M., Hastie R., and Park, B. (1993). Use of mental frequency distributions to represent variability among members of social categories. *Social Cognition*, **11**(1), 22-43.
- Kuehne, S., Forbus, K., Gentner, D. and Quinn, B.(2000) SEQL: Category learning as progressive abstraction using structure mapping. *Proceedings of Cognitive Science Conference.*

- Kuhl, P. K. (1991). Human adults and human infants show a "perceptual magnet effect" for the prototypes of speech categories, monkeys do not. *Perception and Psychophysics*, **50**, 93-107.
- Leake, D. (Ed.) 1996. Case-based Reasoning: Experiences, Lessons and Future Directions, MIT Press.
- Lynch, K. (1960). The image of the city. Cambridge, MA: MIT and Harvard University Press.
- Malmi, R. A., and Samson, D.J. (1983). Intuitive Averaging of Categorized Numerical Stimuli, *Journal of Verbal Learning and Verbal Behavior*, 22, 547-559.
- Malt, B. and Smith, E. (1984). Correlated Properties in Natural Categories. *Journal of Verbal Learning and Verbal Behavior*, 23(2), 250-269.
- McRae, K. (1992). Correlated properties in artifact and natural kind concepts. In *Proceedings of Cognitive Science Conference*, 349-354.
- Paritosh, P.K. and Forbus, K.D. (2003). Qualitative Modeling and Similarity in Back of the Envelope Reasoning. In *Proceedings of the 25th Cognitive Science Conference.*
- Peterson, C.R., and Beach, L.R. (1967). Man as an intuitive statistician, *Psychological Bulletin*, **68**(1), pp 29-46.
- Rips, L. J., and Turbull, W. (1980) How big is big? Relative and absolute properties in memory. *Cognition*, 8, 145-174.
- Rosch, E. (1975). Cognitive Reference Points. *Cognitive Psychology*, **7**, 532-547.
- Ryalls, B. O. and Smith, L. B. (2000). Adults Acquisition of Novel Dimension Words: Creating a Semantic Congruity Effect, *Journal of General Psychology*, **127**(3), 279-326.
- Sachenbacher, M. and Struss, P. (2000). Automated determination of qualitative distinctions: Theoretical foundations and Practical Results, In 14th International Workshop on Qualitative Reasoning, Morelia, Mexico, 144-153.
- Sethna, J. P. (1992). Order parameters, broken symmetry and topology. In *1991 Lectures in Complex Systems* (Eds. L. Nagel and D. Stein), Santa Fe Institute Studies in Sciences of Complexity, Proc. Vol XV, Addison-Wesley, 1992.
- Shum, M. S. (1998). The role of temporal landmarks in autobiographical memory processes. *Psychological Bulletin.* **124**(3), 423-442.
- Skorstad, J., Gentner, D., and Medin, D. (1988). Abstraction processes during concept learning: A structural view. In Proceedings of the Tenth Annual Conference of the Cognitive Science Society, 419-425.
- Staab, S. and Hahn, U. (1998). Grading on the Fly. In Proceedings of the 20th Annual Meeting of the Cognitive Science Society, Madison, WI.
- Tighe, T.J. and Shepp, B.E. (1983). *Perception, cognition* and development: Interactional analysis. Erlbaum Associates, Hillsdale, NJ.
- Tversky, A., and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases, *Science*, 185, pp 1124-1131.
- Tversky, A. (1977). Features of Similarity, *Psychological Review* 84(4), pp 327 352.