

Representing Magnitude by Memory Resonance

A Hypothesis on Qualitative Judgment

Marten J. den Uyl

*Psychological Laboratory
University of Amsterdam*

ABSTRACT

Qualitative judgment, the ability to evaluate attributes that imply some degree of 'goodness' or preference, poses important problems for the information processing paradigm. In this paper one form of qualitative judgment, contextual judgment of magnitude, is analyzed in some detail.

The results of psychophysical experiments are consistent with the idea that human magnitude representation is based on a contextual coding process in which an actual stimulus is compared with a sample of traces of previously encountered similar stimuli. Such a coding process is hard to realize in a conventional memory system.

A distributed model for contextual magnitude judgment is described, in which this trace sampling process is feasible, when special provisions for the use of resonance information are made. Resonance coding involves the representation within a memory system of the memory activity caused by specific patterns of stimulation. A possible implementation of resonance coding, detection of dissonance, is briefly described. The hypothesis is put forward that evaluation of memory resonance plays an equally important role in other forms of qualitative judgment.

INTRODUCTION

People routinely make qualitative judgments: 'an interesting novel', 'a very elegant solution', 'a hot bath', 'a bitter disappointment', 'a highly complex problem'. Such judgments have in common that on the one hand some *classification* of a quality is indicated -interestingness, elegance, hotness, etc.- and on the other hand there is an *evaluation of the degree*, of the quality present.

Qualitative judgment poses pervasive problems to the information processing paradigm. One set of problems is apparent in the actual performance of presentday AI (expert) systems. It is generally recognized that AI systems show distinct patterns of strengths and weaknesses when compared to standards of human intelligence. Where AI systems may easily excel in computational power in formal domains, in 'verbatim' memory capacity, in deductive inference, it is a challenging task to approach human standards in heuristic inference, common sense, creativity, esthetic judgment, and robustness in the face of unexpected and deviant situations. Much of this pattern of weaknesses can be roughly summarized in the statement that AI systems have problems in dealing with *qualitative* aspects of tasks and situations.

Problems with qualitative judgment may be further illustrated in the relative neglect in cognitive psychology of affective processes such as motivation, valuation, emotion (e.g. Mandler 1985) in which qualitative judgment plays a central role.

Yet another set of problems prevail at the epistemological level. When people make qualitative judgments, they experience certain qualitative contents or *qualia*, a tomato 'looks bright red', we 'feel a pain' in a sore tooth. Qualia play a prominent role in debates on the philosophy of mind, and particularly in arguments on the adequacy of the functional (cognitive) paradigm in psychology. It has often been argued that the information processing paradigm cannot account for the phenomenal qualities of human experience.

The major aim of this paper is to outline and illustrate a global hypothesis on the nature of qualitative judgment. According to this hypothesis qualitative judgment is based on an evaluation of *memory resonance*, roughly the impact of stimulation on patterns of autonomous activity in the representational system. The body of the paper is devoted to an analysis of the *archetype* of qualitative judgment: the contextual judgment of magnitude.

QUALITATIVE JUDGMENT OF MAGNITUDE

Any system that operates in the real world must face the elementary task of representing the magnitudes or intensities of objects on continuous physical variables such as length, weight, sound- or light intensity. The question how human beings deal with this task is the proper domain of psychophysics, one of the oldest research traditions in modern experimental psychology. Research in psychophysics has demonstrated many peculiarities of human magnitude representation (For reviews see e.g. Carterette & Friedman 1974).

First of all, magnitude judgment often is a hard task for human beings. The point is perhaps appreciated best when it is realized how dependent we are on the many measurement instruments -from yardstick and balance to sophisticated electronic devices- that have been developed to overcome the limitations of sensory systems.

Problems with magnitude judgment arise in particular when memory representations of magnitudes are involved. Human capabilities in intensity resolution *per se* -i.e. discriminating between magnitudes when two or more are present simultaneously or in immediate succession- differ widely between sensory modalities, but in some cases human discrimination is fair to any standards and may be hard to equal by mechanical devices. One of the few general observations that can be made about intensity discrimination is known as 'Weber's Law': just notable differences (jnd's) between intensities tend to be constant fractions of stimulus intensity.

However, when a single stimulus at a time is judged in relation to a memory representation of one or more reference magnitudes, human performance is severely limited and, moreover, shows surprisingly little variation over sensory modalities. As Miller (1956) noticed, for many different sensory continua people cannot reliably distinguish more than 7 +/- 2 levels of intensity in identification tasks. Performance in intensity identification tasks, where reference levels must be remembered, contrasts sharply with performance in intensity discrimination tasks. For continua like loudness and brightness subjects may be able to partition the range from the weakest, just perceptible, to the largest, just bearable, intensity (the *dynamic range*) in well over a hundred jnd steps, when stimuli are pairwise compared; yet Ss can at most distinguish about 10 levels of intensity with single stimulus presentations.

It may be conjectured that these memory limitations are a direct consequence of the characteristics of the hardware in which human intelligence is implemented. Artificial recording devices store information in physically stable structures that may remain essentially unchanged over longer periods of time. The human brain is a soft kind of hardware, as living tissue grows, changes and decays over time. Human magnitude representation may be understood as a striving for constancy in an inherently unstable system.

A further characteristic of memory representations of magnitude is their dependency on context, particularly on the range of magnitudes over which judgments are made. The stability of memory representations of magnitude decreases with increasing judgmental range. Two intensity levels that can be easily distinguished when they form the extreme magnitudes in a small range stimulus set, may become highly confusable in a large range stimulus set.

Context dependency is also a striking feature of the way magnitudes are typically expressed in everyday language. Magnitudes are typically communicated in qualitative terms like 'big', 'small', 'many', 'few'. Such terms express magnitude in reference to a contextually determined norm; 'a lot of people' implies different numbers at cocktail parties or mass meetings.

How are contextual norms represented in memory? How are judgments made in relation to these norms? One approach, exemplified in the influential 'Adaptation Level' theory (Helson 1964), is to assume that judgments are made in relation to some neutral point that represents the central tendency of the stimulus distribution. Context effects then are represented in a single parameter, the adaptation level. However, it has repeatedly been demonstrated that not just the central tendency but the entire shape of the stimulus distribution is reflected in magnitude judgments (e.g. Parducci & Perrett 1971). Apparently people employ some 'multi parameter' representation of stimulus distributions.

Trace Sample Theory

The characteristics of human magnitude representation just described can be accounted for by a general model of qualitative magnitude judgment ('Trace Sample' theory, den Uyl 1981). Only a few elements of this theory need to be mentioned here.

Magnitude information may be represented in one of two different formats: primary or *trace code* and secondary or *contextual code* (cf. Durlach & Braida 1969). A trace code, the output of some perceptual system, is a direct or analogue representation of stimulus intensity. The precise nature of trace codes we leave unspecified for the moment; it is, however, assumed that trace codes are highly unstable and decay rapidly over time. Hence, the long term memory representation of individual trace codes will be subject to very large error.

The central assumption in trace sample theory is that contextual magnitude representations are formed by comparing the trace code for an actual stimulus magnitude to a sample of trace codes of previously encountered similar stimuli. The resulting contextual magnitude judgment is essentially the rank or percentile score of the actual stimulus magnitude in a subjective reference distribution formed by the trace sample. Thus, when we judge a dog to be 'very large', this essentially means that this dog appears to be larger than most dogs we have seen before. It should be noted that because of the large error in the memory

representation of traces, the subjective reference distribution will be systematically distorted with regard to the objective distribution of reference magnitudes.

Space does not permit to review the evidence here that this simple contextual coding mechanism can account for many findings in experimental psychophysics (cf. den Uyl 1981). We should turn attention to the implications of this coding scheme for the organization of memory.

Classification, Judgment and Memory

Trace Sample theory implies that quite a number of specific traces must be somehow represented in memory; in order to make qualitative judgments a sample of traces should be available for each dimension of each category in memory.

How is the representation of these traces integrated in the representation of categories in memory? One possibility is to assume that *fixed sets of traces* are stored with each category in memory, and can be accessed when the entry for the category is reached as a result of some classification process. The problem with this proposal is that it may easily lead to a proliferation of postulated trace samples when the aim is to account for the flexibility of human judgment. To just mention some potential problems:

-Often the norm for one dimension is dependent on the value on some other dimension, e.g. a height of three feet would be 'tall' for a two year old child, 'very tall' for a 1.5 year old, 'extremely tall' for a one year old. It may be possible to represent dependencies between correlated continuous variables in sets of discrete subcategories (e.g. Lebowitz 1985), but it would not seem a particularly elegant solution.

-Sometimes contextual judgments appear to be made in reference to *ad hoc categories* (Barsalou 1983) constructed on the spot. For example, the judgment that there are 'not so many people present' at a particular lecture, may be made in reference to an ad hoc norm for 'this kind of lecture', defined by a set of circumstances like size of the lecture hall, fame of the lecturer, time of day etc.

An alternative approach that could solve these problems in a principled way, would be to assume that *a trace sample is newly composed for each occasion* where a qualitative judgment is made. That is, for each judgment the memory system samples the traces from the previously judged objects *most similar* to the object presently being judged. Clearly, this is an attractive possibility in that it could conceivably give a system the flexibility apparent in human judgment¹. However, the computational costs of this scheme would seem extravagant in a conventional computational architecture: the scheme implies that some similarity metric is computed between the judged object and each individual object stored in memory. In order to further explore this proposal, we need to consider an implementation of the trace sample model in a parallel-distributed processing architecture in which the implied computations are feasible.

¹An elegant treatment of 'norm theory' based on related principles has recently been presented in (Kahneman & Miller 1986).

MAGNITUDE REPRESENTATION IN A DISTRIBUTED MEMORY

The distributed model I will outline here is based on the 'Harmonium' model developed by Smolensky & Riley (1984) with a few modifications inspired by related models (e.g. McClelland & Rumelhart 1985). Harmony theory relies on a formal mapping between parallel computation and thermal physics and is similar in this respect to the 'Boltzman machine' described by Ackley, Hinton & Sejnowski (1985).

A distributed memory consists of a -large- set of interconnected modules, each of which in turn consists of many interconnected simple processing units or 'nodes'. The activation state of nodes may vary, and the basic mode of operation of nodes is the passing of activation signals to other nodes within and between modules. Specific patterns of activity over the nodes in the memory network constitute active knowledge states. Information processing takes the form of chains of knowledge states, brought about by the units spreading activation through the network.

Each module contains two layers of nodes: *representation nodes* (R-nodes) define the active knowledge states of the system; *trace nodes* (T-nodes) contain information on past contingencies between R-nodes and may send activation signals to R-nodes on this basis. Connections are only between layers, nodes within a single layer are not connected. Both representation and trace nodes take only two activation values: nodes are either on/active or off/inactive.

The basic cognitive operation in a connectionist memory is pattern completion. Pattern completion takes place in (asynchronous) processing cycles in individual modules as follows: At the beginning of the cycle some subset of the R-nodes in the module is clamped into activation states by incoming connections from other modules. Other R-nodes are assumed to have random activation values at the beginning of the cycle. The task for the module is to reinstate a stable and complete pattern over all the representation nodes in the module. Note that in a distributed model 'psychological' stimulus features (e.g. shapes, color) do not correspond to individual processing units, but to activation patterns over collections of units.

Each trace node is connected to a set of representation nodes by bidirectional links. A T-node contains -as a result of past experience- a *key*, a set of weights on connections with R-nodes with values of either +1 or -1. This key defines a 'preferred pattern' of activation states over connected R-nodes.

A processing cycle is divided in discrete 'ticks' (McClelland & Rumelhart 1985). At each tick a T-node receives an activation signal from each R-node that is consistent with the key (the state of the R-node matches the sign of the weight) and a de-activation signal for each mismatch. A trace node is in the active state as long as the sum of the (de-)activation signals exceeds some variable threshold value. Active T-nodes send activation signals to R-nodes in accordance with the weights in the key, e.g. a de-activation signal is send when the weight is negative.

The thresholds of T-nodes are gradually raised in the course of a processing cycle. As a result many T-nodes may participate in the first stages of a cycle. At the end of a cycle only the best matching T-nodes remain active when the module 'freezes' into a stable completion.

The activation state of a R-node not clamped by external connections is a (probabilistic) function of the incoming signals from T-nodes; representation nodes will tend to conform to the 'majority vote' of incoming signals.

How do trace nodes come to represent contingencies between representation nodes? I will shortcut this complex issue here by postulating a global learning rule that suffices for present purposes: at the end of a completion cycle, key weights of then active trace nodes that are inconsistent with the activation state of the connected representation node may change their sign with a certain probability so as to achieve consistency².

Harmony

In a 'harmonium model' processing is driven by a single principle: that of achieving completions with the highest harmony or 'self-consistency'. Harmony is the degree of consistency between a pattern of activation over the representation nodes, and a set of preferred patterns defined by the keys in the active trace nodes. More precisely, the harmony of a system state can be expressed as follows³:

Let a vector T represent the states of trace nodes with the values active=1 and inactive=0; a vector R over representation nodes takes the values on=+1 and off=-1; a vector K_i represents the key in trace node T_i with the values +1 and -1 for positive and negative weights, and 0 when T_i is not connected to R_j . The harmony of a state then is:

$$H_{(T,R)} = \sum_i T_i R_i K_i \quad (1)$$

It can readily be seen that (1) is equivalent to the sum of all consistent signals in the module minus the summed inconsistent signals ('*' denotes the dot product).

Primary Coding of Magnitude

Magnitudes on some dimension are represented on a subset of the R-nodes in a harmonium module, the 'magnitude representation nodes' or M-nodes. The primary code for a magnitude is simply the number of M-nodes active in the module. This primary coding may be brought about thus.

M-nodes are clamped into activation states by external connections carrying activation signals that have their origin in sensory systems. For each M-node a threshold parameter m_{it} is defined that randomly fluctuates over time. An M-node will be clamped in the active state upon presentation of a stimulus magnitude s_j when $s_j > m_{it}$ and will be clamped 'off' otherwise.

Weber's Law implies that the thresholds of M-nodes are spaced approximately geometrically over the dynamic range of the magnitude dimension. The cumulative distribution function, i.e. the expected number of active M-nodes, then is a logarithmic function of stimulus intensity.

² This postulate marks a transition from an 'enumeration of specific instances' principle in Smolensky & Riley (1984) to a 'superposition of traces' principle (McClelland & Rumelhart 1985).

³ The present notation is slightly different from Smolensky & Riley (1984).

A rationale for the present instantiation of Fechner's time-honored 'Logarithmic Law' is that this coding scheme makes minimal demands on the stability of the units in which magnitude information is encoded. In this way reliable intensity resolution can be achieved by parallelling highly unreliable individual units.

Contextual Coding

In Trace Sample theory (den Uyl 1981) the contextual code for a stimulus magnitude s_j is the proportion of traces in the trace sample smaller than s_j . The present distributed model does not store traces (copies) of magnitudes. Stimuli may only have a lasting impact on memory by changing key weights. Yet, a contextual code can be computed in the harmonium model that is equivalent to the secondary code in Trace Sample theory.

Four kinds of *trace signals* $t(R,K)$ from T-nodes to R-nodes contribute to system harmony:

- activation signals to active R-nodes: $t(+,+)$,
- de-activation signals to active R-nodes: $t(+,-)$,
- activation signals to inactive R-nodes: $t(-,+)$,
- de-activation signals to inactive R-nodes: $t(-,-)$.

Inconsistent trace signals ($t(+,-)$ and $t(-,+)$) provide information concerning the position of s_j in the magnitude distribution of similar stimuli. A signal $t(+,-)$ implies that the present stimulus dominates a threshold ($s_j > m_{it}$) while some past stimulus s_p which caused the negative key weight in the signal, had been below this threshold ($m_{it} > s_p$). Hence, it may be inferred from $t(+,-)$ that s_j dominates at least one past stimulus. Analogously, a signal $t(-,+)$ indicates that s_j is smaller than some past stimulus. Of course, error is introduced into these inferences because of the random fluctuations over time of the thresholds.

Consistent signals do not provide contextual information, these signals only indicate where both s_j and past stimuli stand with regard to threshold nodes.

We denote the frequency of the signal $t(R,K)$ summed over M-nodes, given activation vectors T and M upon presentation of stimulus s_j by $f_{(T,M)j}(R,K)$. A contextual code C_j equivalent to the percentile code in Trace Sample theory can then be expressed as:

$$C_j = \frac{f_{(T,M)j}(+,-)}{f_{(T,M)j}(+,-) + f_{(T,M)j}(-,+)} \quad (2)$$

C_j may range from 0 for extremely small stimuli to 1 for extremely large stimuli, the expected value for intermediate stimuli (the adaptation level) is 0.5.

It may be observed that in the present model norms for magnitude judgment, i.e. the distribution of trace signals in a harmonium module, are indeed composed anew for each judgment, as each stimulus pattern may activate a different T vector and hence sample a different set of key weights (cf. McClelland & Rumelhart 1985).

Dissonance as a Contextual Code

The contextual coding scheme in the preceding section raises one important problem: How can the contextual code in (2) be computed in a distributed system? All the information required for the computation is present in the M-nodes. It would appear a simple solution to have all M-nodes send counts of inconsistent incoming signals to special 'collector-units' where the necessary computations could be performed. However, this proposal would require M-nodes to spread other -more complex- information through the system than just their activation state. The proposal implies a major breach with connectionist design principles.

There are ways to approximate the contextual code using standard -though specialized- connectionist processors. One such scheme involves the detection of *dissonance* in activation patterns. A clamped representation node is dissonant, when a majority of incoming trace signals are inconsistent with the clamped activation state of the R-node.

Dissonance can be detected in the following way. Suppose a completion cycle in a harmonium module is followed by a *resonance cycle* in which the activation pattern on T-nodes remains unchanged, but all clamps on R-nodes are removed, thus allowing R-nodes to settle into preferred states. Dissonant nodes -i.e. R-nodes which change their activation state in the resonance cycle- could then be registered in specialized units (e.g. units sensitive to time-contrast in activation signals).

The contextual code in (2) can be approximated in various ways from the dissonance in M. Large stimuli will tend to clamp high-threshold M-nodes -expected to be off- in dissonant active states, small stimuli will produce dissonance in inactive nodes. Intermediate magnitudes will cause the lowest overall dissonance in M.

The precise form a dissonance-based contextual code may take is not important for the moment. The general point I want to make here is that although it is feasible to develop fairly simple coding schemes for contextual magnitude representations in a connectionist memory module, some special provisions are required to this end.

RESONANCE AND QUALITATIVE JUDGMENT

It may be conjectured that all qualitative judgments share the structure of contextual magnitude judgment, they all involve the contextual evaluation of *memory resonance* to an object description.

A crucial step must be taken in order to use resonance information in a memory system. Global characteristics of memory activity in response to a pattern of stimulation must be represented within the memory system itself, in order to be interpreted as providing information about the object that gave rise to this memory activity. For instance, in order to use the extent of inconsistent activation in a module as a measure of the extremeness of the magnitude of an object, inconsistent activation must in some way be detected and registered.

It has been proposed that resonance evaluation takes the form of evaluating patterns of dissonance between a 'preferred' or expected resonance pattern, and the pattern externally imposed on the memory. A rationale for this form of evaluation can be found in the design principles of a distributed memory. In a distributed memory knowledge is hidden rather than

stored. Past events only leave traces in the form of changes in activation weights. It is not easy to systematically search or recover past episodes from such traces.

Some tasks, like magnitude judgment or evaluation of prototypicality, require a system to make use of knowledge concerning the *distribution of different stimulus patterns* in the past. Yet, a module in a distributed memory can only support one pattern of activation at a time. Traces of past events that do not correspond to an actual pattern in a module may only influence further processing, if it is in some way registered that they would have liked to see things different.

I have said little about 'goodness', preference and affective qualities in qualitative judgment. The reason is that liking does not occur -except in a metaphorical sense- on the level of memory modules. Only the system as a whole has likes and dislikes. In order to extend the present hypothesis to affective evaluation, the notion of memory resonance should be extended to higher levels of organization, i.e. resonance should be evaluated over collections of modules rather than within single modules. A suitable framework for such an extension can be found in Frijda's 'concern-realization' theory of emotion (Frijda 1986; den Uyl & Frijda 1984).

REFERENCES

- Ackley, D., Hinton, G. E. & Sejnowski, T. J. (1985). Boltzmann machines: Constraint satisfaction networks that learn. *Cognitive Science* 9, 147-169.
- Barsalou, L. W. (1983) Ad hoc categories. *Memory & Cognition*, 11, 211-227.
- Carterette, E. C. & Friedman, M. P. (1974). *Handbook of Perception Volume II*, Ac. Press, New York.
- Durlach, N. I. & Braida, L. D. (1969). Intensity Perception I: Preliminary theory of intensity resolution. *Journal of the Acoustical Society of America*, 46, 372-383.
- Frijda, N.H. (1986) *The Emotions*. New York: Cambridge U.P.
- Helson, H. (1964). *Adaptation Level Theory*, New York, Harper and Rowe.
- Kahneman, D. & Miller, D.T. (1986). Norm Theory: comparing reality to its alternatives. *Psychological Review*, Vol. 93, 136-153.
- Lebowitz, M. (1985). Categorizing numeric information for generalization. *Cognitive Science*, 9, 285-308
- Mandler, G. (1985). *Cognitive Psychology: an essay in cognitive science*. LEA, Hillsdale, N.J.
- McClelland, J. L. & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, Vol. 114, 159-188.
- Miller, G. A. (1956). The magical number 7 plus or minus two. *Psychological Review*, 63, 81-97.
- Parducci, A. & Perrett, L.F. (1971) Category rating scales: effects of relative spacing and frequency of stimulus values. *Journal of Experimental Psychology, Monograph*, 89, 427-452.
- Smolensky, P. & Riley, M. S. (1984). Harmony Theory: Problem solving, parallel cognitive models, and thermal physics. ICS Report 8404, UC San Diego.
- den Uyl, M.J. (1981). On the Relativity of Human Judgment: Outline of a Theory of Contextual Magnitude Representation. Manuscript, Psych. Dept. Stanford University.
- den Uyl M.J. & Frijda, N.H. (1984). Mood, Emotion and Action: A Concern-Realization Model. Proceedings of the Sixth Cognitive Science Conference, Boulder, Co. 137-141.