

AD-A137 883

HARMONY THEORY: A MATHEMATICAL FRAMEWORK FOR STOCHASTIC 1/1

PARALLEL PROCESSING(U) CALIFORNIA UNIV SAN DIEGO LA
JOLLA INST FOR COGNITIVE SCIENCE P SMOLENSKY DEC 83

UNCLASSIFIED

ICS-8306 N00014-79-C-0323

F/G 12/1

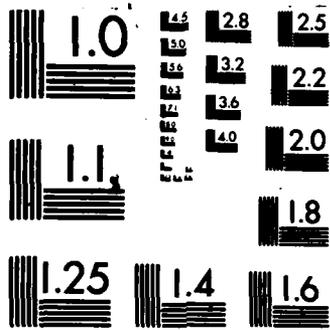
NL

END

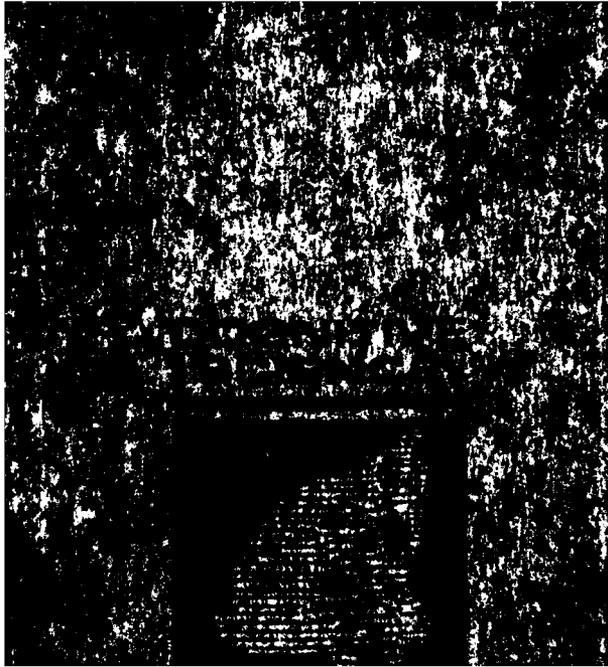
FILMED

3-

DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A



Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

1

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ONR 8304	2. GOVT ACCESSION NO. AD-A137883	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Harmony Theory: A Mathematical Framework for Stochastic Parallel Processing		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Paul Smolensky		8. CONTRACT OR GRANT NUMBER(s) N00014-79-C-0323
9. PERFORMING ORGANIZATION NAME AND ADDRESS Center for Human Information Processing Institute for Cognitive Science University of California, San Diego La Jolla, Calif. 92093		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 667-437
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training (Code 442) Office of Naval Research 800 No. Quincy St., Arlington, VA 22217		12. REPORT DATE December, 1983
		13. NUMBER OF PAGES 9
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) This research was also supported by a grant from the System Development Foundation, as well as a grant from the Alfred P. Sloan Foundation during the 1981-82 academic year.		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) stochastic parallel computation temperature parallel processing schema theory statistical machines learning neural modelling perception		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		

FEB 15 1984

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102 LF 014 6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

**HARMONY THEORY:
A MATHEMATICAL FRAMEWORK FOR
STOCHASTIC PARALLEL PROCESSING**

Paul Smolensky
Institute for Cognitive Science
University of California, San Diego C-015
La Jolla, CA 92093

Published in
Proceedings of the National Conference on Artificial Intelligence
AAAI-83; Washington, DC; August, 1983

I am indebted to George Mandler and Donald Norman for the opportunity to pursue this work. Geoffrey Hinton, James McClelland, David Rumelhart, and the members of the UCSD Parallel Distributed Processing research group have contributed enormously to the perspective on cognition taken in this work. I am grateful to Francis Crick, Mary Ann McCrary, Michael Mower, and Donald Norman for extremely helpful comments on the paper. Above all, warm thanks go to Douglas Hofstadter for his continual encouragement; the influence of his ideas pervades this research.

The research reported here was conducted under Contract N00014-79-C-0323, NR 667-437 with the Personnel and Training Research Programs of the Office of Naval Research, and was sponsored by the Office of Naval Research and a grant from the System Development Foundation. During the 1981-82 academic year, support was provided by the Alfred P. Sloan Foundation and Grant PHS MH 14268 to the Center for Human Information Processing from the National Institute of Mental Health. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsoring agencies. Approved for public release; distribution unlimited. Reproduction in whole or in part is permitted for any purpose of the United States Government.

Introduction

Our sensory systems are capable of representing a vast number of possible stimuli. Our environment presents us with only a small fraction of the possibilities; this selected subset is characterized by many regularities. Our minds encode these regularities, and this gives us some ability to infer the probable current condition of unknown portions of the environment given some limited information about the current state. What kind of regularities exist in the environment and how should they be encoded?

This paper presents preliminary results of research founded on the hypothesis that in real environments there exist regularities that can be idealized as mathematical structures that are simple enough to be analyzable. Only the simplest kind of regularity is considered here: I will assume that the environment contains modules (objects) that recur exactly, with various states of the environment being comprised of various combinations of these modules. Even this simplest kind of environmental regularity offers interesting learning problems and results. It also serves to introduce a general framework capable of treating more subtle types of regularities. And the problem considered is an important one, for the delineation of modules at one level of conceptual representation is a major step in the construction of higher level representations.

To analyze the encoding of modularity of the environment, I will proceed in three steps. First, I will describe a general information processing task, *completion*, for a cognitive system. Then I will describe the entities, *schemas*, I use for encoding the environmental modules and discuss how they are used by the cognitive system to perform its task. Finally I will offer a criterion for how the encoding of the environment into schemas should be done. The presentation will be informal; more precise statements of definitions and results appear in the Appendix.

Completions and Schemas

The task I consider here, *completion*, is the inferring of missing information about the current state of the environment. The cognitive system is given a partial description of that state as input, and must produce as output a completion of that description.

The entities used to represent the environmental modules are called *schemas*. A given schema represents the hypothesis that a given module is present in the current state of the environment. When the system is given a partial description of the current state of the environment, the schemas look at the information to see if they fit it; those that do become *active*. When inference from the given information to the missing information is possible, it is because some of the schemas represent modules that incorporate both given and unknown information. Such a schema being active, i.e., the belief that the module is present in the current environmental state, permits inferences about the missing information pertaining to the module. (Thus if the given information about a word is ALG###THEM, the schema for the module *algorithm* is activated, and inferences about the missing letters are possible.)

There is a problem here with the sequence of decision-making. How can the schemas decide if they fit the current situation when that situation is only partially described? It would appear that to really assess the relevance of a schema, the system would first have to fill in the missing information. But to decide on the missing portions, the system first needs to formulate

beliefs concerning which modules are present in the current state. I call this the *schema / inference decision problem*; what is desired is a way of circumventing this problem that is general and extremely simple to describe mathematically.

Harmony and Computational Temperature

Before discussing the algorithm used to tackle this problem, let us first try to characterize what outcome we would like the algorithm to produce. The approach I am pursuing assumes that *the best inference about the missing information is the one that best satisfies the most hypotheses (schemas)*. That is, consider all possible responses of the system to an input (a partially specified environmental state). Each such response involves (a) a decision about exactly which schemas are active, and (b) a decision about how to specify all the missing information. Given such a response, I propose a measure of the *internal consistency* that computes the degree to which the hypotheses represented by the active schemas are satisfied by the input and output. I call this measure the *harmony function* H . Responses characterized by greater internal consistency — greater harmony — are "better." (The definition of H is simple, but requires the formalism given in the Appendix.)

Given that better responses are characterized by greater harmony, the obvious thing to do is hill-climb in harmony. However this method is apt to get stuck at local maxima that are not global maxima. We therefore temporarily relax our desire to go directly for the "best" response. Instead we consider a *stochastic* system that gives different responses with different probabilities: the better the response, the more likely the system is to give it. The degree of spread in the probability distribution is denoted T ; for high values of T , the distribution is widely spread, with the better responses being only slightly more likely than less good ones; for low values of T , the best response is much more likely than the others. Thus T measures the "randomness" in the system; I call it the *computational temperature*. When we want only good responses, we must achieve low temperature.

A general stochastic algorithm can be derived that realizes this probabilistic response; it provides a method for computer simulation. A parallel relaxation method is used to resolve the *schema / inference decision problem* in a way that involves no sophisticated control.

The variables of the system are the activations of all the schemas (1 = active, 0 = inactive) and the bits of missing information. The system starts with (a) all schemas inactive, and (b) completely random guesses for all bits of missing information. Then, a schema or a bit of missing information is selected randomly as the variable to be inspected; it will now be assigned a (possibly new) value of 1 or 0. Using the current guesses for all the other variables, the harmony function is evaluated for the two possible states of the selected variable. These two numbers, $H(1)$ and $H(0)$, measure the overall consistency for the two cases where the selected variable is assigned 1 and 0; they can be computed because tentative guesses have been made for all the schema activations and missing information. Next a random choice of these 1,0 states is made, with the probability of the choices 1 and 0 having ratio $\frac{\text{prob}(1)}{\text{prob}(0)} = \frac{e^{H(1)/T}}{e^{H(0)/T}}$. (The reason for the exponential function is given in the Appendix.)

This random process — pick a schema or bit of missing information; evaluate the two harmonies; pick a state — is iterated. In the theoretical limit that the process continues indefinitely, it can be proved that the probability that the system gives any response is

proportional to $e^{H/\pi}$, where H is the harmony of that response. This probability distribution satisfies the qualitative description of response probabilities we set out to realize.

Cooling the System

In this algorithm, each schema activation or bit of missing information is determined randomly; most likely the value with higher harmony is chosen, but sometimes not. Thus most of the time the changes raise the harmony, but not always. The higher the temperature, the more random are the decisions, that is, the more often the changes go "downhill." Thus, this algorithm, unlike strict hill-climbing, does *not* get stuck at local maxima. However, there is a tradeoff. The higher the temperature, the faster the system escapes local maxima by going downhill; but the higher the temperature, the more random is the motion, and the more time the system spends in states of low harmony. Eventually, to be quite sure the system's response has high harmony, the temperature must be low.

As the computation proceeds, the optimization point of this tradeoff shifts. Initially, the guesses for the missing information are completely random, so the information on which schemas are determining their relevance is unreliable. It is therefore desirable to have considerable randomness in the decision making, i.e., high temperature. Even at high temperature, however, the system is more likely to occupy states of higher harmony than lower, so the guesses become more reliable than their completely random start. At this point it makes sense for the schemas to be somewhat less random in their activity decisions, so the temperature should be lowered a bit. This causes the system to spend more of its time in states of higher harmony, justifying a further decrease in temperature. And so the temperature should be gradually lowered to achieve the desired final condition of low temperature. (A major portion of future work will be analysis of how to regulate the cooling of the system.)

As the computation proceeds and the temperature drops, the system's initially rough and scattered response becomes progressively more accurate and consistent. This is just the kind of computation typical in people and just the kind needed in any large parallel system, where each sub-system needs a constant stream of input from the others.

Cognitive Crystallization

As computation proceeds, does accuracy increase slowly and steadily, or does the system undergo sudden and dramatic changes in behavior, as do physical systems when they are cooled past critical temperatures marking phase transitions? This question has been addressed both through computer simulation and analytic approximation of a two-choice decision. The system has two schemas representing conflicting interpretations of the environment. The approximate theory allows computation of the probabilities of various completions given an input that partly describes the environment. It is useful to ask, *what is the completion of a completely ambiguous input?* For high temperatures, as one might expect, the completions form random mixtures of the two interpretations; the system does not choose either interpretation. However, *below a certain "freezing temperature," the system adopts one of the interpretations.* Each interpretation is equally likely to be selected. The computer simulation approximates this behavior; below the freezing point, it flips back and forth between the two interpretations, occupying each for a long time.

Slowly vacillating interpretation of genuinely ambiguous input is a familiar but not particularly important feature of human cognition. What is significant here is that *even when the input provides no help whatever in selecting an interpretation, the system eventually (when cooled sufficiently) abandons meaningless mixtures of interpretations and adopts some coherent interpretation.* A robust tendency to form coherent interpretations is important both for modelling human cognition and for building intelligent machines. The above analysis suggests that in processing typical inputs, which are at most partially ambiguous, as processing continues and the temperature drops, the system wanders randomly through an ever-narrowing range of approximate solutions until some time when the system freezes into an answer.

Schema Selection

Having discussed how schemas are used to do completions, it is time to consider what set of schemas ought to be used to represent the regularities in a given environment. Suppose the system experiences a set of states of the environment and from these it must choose its schemas. Call these states the *training set*. Since the schemas are used to try to construct high-harmony responses, a reasonable criterion would seem to be: *the best schemas are those that permit the greatest total harmony for responses to the training set.*

I call this the *training harmony criterion*. I will not discuss here an algorithm for finding the best schemas. Instead, I shall present some elementary but non-obvious implications of this very simple criterion. To explore these implications, we create various idealized environments each displaying some interesting modularity. Choosing a training set from within an environment, we see whether the training harmony criterion allows the system to *induce the modularity from the training set*, by choosing schemas that encode the modularity.

Perceptual Grouping

We perceive scenes not as wholes, nor as vast collections of visual features, but as collections of objects. Is there some general characterization of what is natural about the particular levels of grouping that form these "objects"? This question can be addressed at a simple but abstract level by considering an environment of strings of four letters in some fixed font. In this idealized environment, the modules ("objects") are letter tokens, and in various states of the environment they recur exactly: each letter always appears in exactly the same form. The location of each of the four letters is absolutely fixed; I call the location of the first letter the "first slot," and so forth. The environment consists of all combinations of four letters.

Now consider a computer given a subset of these four-letter strings as training; call these the *training words*. The image the computer gets of each training word is just a string of bits, each bit representing whether some portion of the image is on or off. The machine does not know that "bit 42" and "bit 67" represent adjacent places; all bits are spatially meaningless. Could the machine possibly induce from the training that certain bits "go together" in determining one of the letters, for example, those bits that we know represent the first slot? Is there some sense in which schemas for the letter A in the first slot, and so on, are natural encodings of this environment?

As an obvious alternative, for example, the system could simply create one schema for each training word. Or it could create a schema for each bit in the image. These are the two extreme cases of maximally big and small schemas; the letter schemas fall somewhere in between. Which of these three cases is best? *The training harmony criterion implies that letter schemas are best*, provided the training set and number of bits per letter are not too small.

This result can be abstracted from the reading context in which it was presented for expository convenience; the mathematical result does not depend upon the interpretation we place upon the modules with which it deals. Thus the result can be characterized more abstractly as providing evidence that *natural schemas encoding the modules of an environment are inducible by the training harmony criterion*, provided the modules recur exactly. This investigation must now be extended to cases in which the recurrence of the modules is in some sense approximate.

Rules vs. Instances

When should experience be encoded as a list of instances and when as a collection of rules? To address this issue we consider two environments that are special subsets of the four-letter environment considered above:

Environment <i>R</i>				Environment <i>I</i>			
FAMB	VEMB	SIMB	ZOMB	FARB	VAMP	SALT	ZAND
FAND	VEND	SIND	ZOND	FENP	VELB	SEMD	ZERT
FARP	VERP	SIRP	ZORP	FIMT	VIRD	SINB	ZILP
FALT	VELT	SILT	ZOLT	FOLD	VONT	SORP	ZOMB

In the highly regular environment *R*, there are strict rules such as "F is always followed by A"; in the irregular environment *I*, no such rules exist. Note that here, schemas for the "rules" of environment *R* are just *digraph* schemas:

FA—, *VE*—, . . . , —*RP*, —*LT*;

schemas for "instances" are whole-word schemas.

One might hope that a criterion for schema selection would dictate that environment *R* be encoded in digraph schemas representing the rules while environment *I* be encoded in word schemas representing instances. *The training harmony criterion implies that for the regular environment, digraph schemas are better than word schemas; for the irregular environment, it is the reverse.* (In each case the entire environment is taken as the training set.)

Higher Level Analyses

The framework described here is capable of addressing more sophisticated learning issues. In particular, it is well-suited to analyzing the construction of higher-level representations and considering, for example, the value of hierarchical organization (which is not put into the system, but may come out in appropriate environments). In addition to addressing issues of schema selection at the more perceptual levels considered here, the framework can be employed at higher conceptual levels. The selection of temporal scripts, for example, can be considered by taking the environment to be a collection of temporally extended episodes. The simulation method described here for systems with given schemas can also be applied at higher levels; it is being explored, for example, for use in text comprehension.

APPENDIX: The Formal Framework of Harmony Theory

In the following general discussion, the specifics of the four-letter grouping environment discussed in the text are presented in parentheses.

The possible beliefs B of the cognitive system about the current state of the environment form a space B . This space is assumed to have a set p of binary coordinates p ; every belief B is defined by a set of bits $B(p)$ in $\{+1, -1\}$, one for each $p \in p$. (Each p represents a distinct pixel. $B(p)$ is $+1$ if the system believes p is on, -1 if off. This belief can come from inference or input.) An input I to the system is a set of binary values $I(p)$ for some of the $p \in p$; an output is a set of binary values $O(p)$ for the remaining p . Together, I and O form a complete belief state B , a completion of I .

A schema S is defined by a value $S(p) \in \{+1, -1, 0\}$ for each $p \in p$. (If pixel p does not lie in the first slot, the schema S for "the letter A in the first slot" has $S(p) = 0$. If p does lie in the first slot, then $S(p)$ is ± 1 according to whether the pixel is on or off.) If for some particular p , $S(p) \neq 0$, then that p is an argument of S ; the number of arguments of S is denoted $|S|$. (For a word schema, every p is an argument.)

Schemas are used by the system to infer likely states of the environment. For a given environment, some schemas should be absent, others present, with possibly varying relative strengths corresponding to varying likelihood in the environment. A knowledge base for the system is a function σ that defines the relative strengths of all possible schemas: $\sigma(S) \geq 0$ and $\sum_S \sigma(S) = 1$. (The knowledge bases relevant to the text have all $\sigma(S) = 0$ except for the schemas in some set S — like letters — for which the strengths are all equal. These strengths, then, are all $1/|S|$, the inverse of the number of schemas in the set.)

A response of the system to an input I is a pair (A, B) , where B is a completion of I , and A defines the schema activations: $A(S) \in \{0, 1\}$ for each schema S .

A harmony function H is a function that assigns a real number $H_\sigma(A, B)$ to a response (A, B) , given a knowledge base σ , and obeys certain properties to be discussed elsewhere. A particular exemplar is

$$H_\sigma(A, B) = \sum_S \sigma(S) A(S) \left[\sum_p B(p) S(p) - \kappa |S| \right]$$

Here κ is a constant in the interval $[0, 1]$; it regulates what proportion ϵ of a schema's arguments must disagree with the beliefs in order for the harmony resulting from activating that schema to be negative: $\epsilon = \frac{1}{2}(1 - \kappa)$. In the following we assume ϵ to be small and positive. (The terms without κ make H simply the sum over all active schemas of the strength of the schema times the number of bits of belief (pixels) that are consistent with the schema minus the number that are inconsistent. Then each active schema incurs a cost proportional to its number of arguments, where κ is the constant of proportionality.)

The probability of any response (A, B) is a monotonically increasing function f of its harmony $H(A, B)$. f is constrained by the following observations. If p_1 and p_2 are not connected — even indirectly — through the knowledge base σ , then inferences about p_1 should be

statistically independent of those about p_2 . In that case, H is the sum of the harmony contributed by three sets of schemas: those connected to p_1 , those connected to p_2 , and those connected to neither. The desired statistical independence requires that the individual probabilities of responses for p_1 and p_2 multiply together. Thus f must be a function that takes additive harmonies into multiplicative probabilities. The only continuous functions that do this are the exponential functions, a class that can be parametrized by a single parameter T ; thus

$$\text{prob}(A, B) = n e^{H(A, B)/T}.$$

The normalization constant n is chosen so that the probabilities for all possible responses add up to one. T must be positive in order that the probability increase with H . As T approaches zero, the function f approaches the discontinuous function that assigns equal nonzero probability to maximal harmony responses and zero probability to all others.

Considerations similar to those of the previous paragraph lead to an isomorphic formula in thermal physics (the Boltzmann law) relating the probability of a physical state to energy. The randomness parameter there is ordinary temperature, and therefore, I call T the *computational temperature* of the cognitive system.

The probability of a completion B is simply the probability of all possible responses (A, B) that combine the beliefs B with schema activations A : $\text{prob}(B) = n \sum_A e^{H(A, B)/T}$.

Monte Carlo analyses of systems in thermal physics have proven quite successful (Binder, 1976). Starting from the exponential probability distribution given above, the stochastic process described in the text can be derived (as in Smolensky, 1981). The above formula for H leads to a simple form for the stochastic decision algorithm of the text which supports the following interpretation. The variables can be represented as a network of " p nodes" each carrying value $B(p)$, " S nodes" each carrying value $A(S)$, and undirected links from each p to each S carrying label $\sigma(S) \cdot S(p)$ (links labelled zero can be omitted). The nodes represent stochastic processors running in parallel, continually transmitting their values over the links and asynchronously setting their values, each using only the labels on the links attached to it and the values from other nodes transmitted over those links. (The values $I(p)$ of those p nodes determined by the input are fixed.) This representation makes contact with the neurally-inspired "connectionist" approach to cognition and parallel computation (Hinton and Anderson, 1981; McClelland and Rumelhart, 1981; Rumelhart and McClelland, 1982). Independently of the development of harmony theory, Hinton and Sejnowski (1983) developed a closely related approach to stochastic parallel networks following (Hopfield, 1982) and (Kirkpatrick, Gelatt, and Vecchi, 1983). From a nonconnectionist artificial intelligence perspective, Hofstadter (1983) is pursuing a related approach to perceptual grouping; his ideas have been inspirational for my work (Hofstadter, 1979).

An *environment* for the cognitive system is a probability distribution on B . (All patterns of on/off pixels corresponding to sequences of four letters are equally probable; all other patterns have probability zero.) A *training set* T from this environment is a sample of points T drawn from the distribution. (Each T is a training word.) A *response* A of the system to the set T is a specification of schema activations $A(T)$ for each T . The *training harmony* of such a response is $H_\sigma(A, T) = \sum_T H_\sigma(A(T), T)$. The maximum of $H_\sigma(A, T)$ over all responses A is $H_\sigma(T)$, the *training harmony permitted by the knowledge base* σ .

Each of the sets of schemas considered in the text (letters, digraphs, etc.) tile the training set T , in that each T agrees exactly with schemas that have nonoverlapping arguments and that together cover all of T . All results cited in the text follow from this elementary calculation: *If the knowledge base σ consists of a set of schemas S that tile T , then $H_\sigma(T) = \text{const.}/\#S$ where const. is a constant for a given T .* Thus, given a single training set T and two tilings of T with sets of schemas, the set with fewer schemas permits greater training harmony, and is preferred by the training harmony criterion. If the number of letters in the alphabet a is smaller than the number of pixels per letter, letters are a better encoding than pixels; if the number of training words exceeds $4a$ then letters are better than words. The number of word schemas needed in each of the restricted environments R and I is 16; the number of digraphs needed for R is 8 while for I it is 96.

The calculation cited in the previous paragraph is quite simple. Recall that if the proportion of a schema's arguments that disagree with the beliefs exceeds ϵ , then the harmony resulting from activating that schema is negative. Thus if ϵ is chosen small enough (which we assume), then for any given training word T , only those schemas that match exactly can contribute positive harmony. In the response $A(T)$ that maximizes the harmony, therefore, only exactly matching schemas will be active; the others will be inactive, contributing zero harmony. Since the schemas S tile T , for each pixel p in T there is exactly one active schema with p as an argument, and the value $B(p)$ of the pixel is consistent with that schema, so $B(p)S(p) = 1$. Thus,

$$\sum_S A(S) \sum_p [B(p)S(p) - \kappa |S(p)|] = \#p(1-\kappa) = \#p \cdot 2\epsilon.$$

Because $\sigma(S)$ is a constant for all $S \in S$, the harmony $H_\sigma(A(T), T)$ is simply the previous expression times $\sigma(S) = 1/\#S$, or $2\epsilon\#p/\#S$. Since this quantity is identical for all training words T , summing over all T in the training set T just multiplies this by the size of T , giving $2\epsilon\#p\#T/\#S$ as the training harmony permitted by σ :

$$H_\sigma(T) = \text{const.}/\#S$$

where const. = $2\epsilon\#p\#T$, which is constant for a given T .

REFERENCES

- Binder, K. (1976). Monte Carlo investigations of phase transitions and critical phenomena. In C. Domb & M. S. Green (Eds.), *Phase transitions and critical phenomena, Vol. 5b*. New York: Academic Press.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences USA*, 79, 2554-2558.
- Hinton, G. E., & J. A. Anderson (Eds.) (1981). *Parallel models of associative memory*. Hillsdale, NJ: Erlbaum.
- Hinton, G. E., & Sejnowski, T. J. (1983, May). Analyzing cooperative computation. *Proceedings of the Fifth Annual Conference of the Cognitive Science Society*. Rochester, NY.
- Hinton, G. E., & Sejnowski, T. J. (1983, June). Optimal perceptual inference. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Washington, DC.
- Hofstadter, D. R. (1979). *Godel, Escher, Bach: An eternal golden braid*. New York: Basic Books.
- Hofstadter, D. R. (1983, June). The architecture of jumbo. *Proceedings of the Machine Learning Workshop*, Illinois.
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220:4598, 671-680.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception, Part 1: An account of the basic findings. *Psychological Review*, 88, 375-407.
- Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception, Part 2: The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, 89, 60-94.
- Smolensky, P. (1981). *Lattice renormalization of ϕ^4 theory*. Unpublished doctoral dissertation, Indiana University.

Cognitive Science ONR Technical Report List

- 8001. Donald R. Gentner, Jonathan Grudin, and Eileen Conway. *Finger Movements in Transcription Typing*. May 1980.
- 8002. James L. McClelland and David E. Rumelhart. *An Interactive Activation Model of the Effect of Context in Perception: Part I*. May 1980.
- 8003. David E. Rumelhart and James L. McClelland. *An Interactive Activation Model of the Effect of Context in Perception: Part II*. July 1980.
- 8004. Donald A. Norman. *Errors in Human Performance*. August 1980.
- 8005. David E. Rumelhart and Donald A. Norman. *Analogical Processes in Learning*. September 1980.
- 8006. Donald A. Norman and Tim Shallice. *Attention to Action: Willed and Automatic Control of Behavior*. December 1980.
- 8101. David E. Rumelhart. *Understanding Understanding*. January 1981.
- 8102. David E. Rumelhart and Donald A. Norman. *Simulating a Skilled Typist: A Study of Skilled Cognitive-Motor Performance*. May 1981.
- 8103. Donald R. Gentner. *Skilled Finger Movements in Typing*. July 1981.
- 8104. Michael I. Jordan. *The Timing of Endpoints in Movement*. November 1981.
- 8105. Gary Perlman. *Two Papers in Cognitive Engineering: The Design of an Interface to a Programming System and MENUNIX: A Menu-Based Interface to UNIX (User Manual)*. November 1981.
- 8106. Donald A. Norman and Diane Fisher. *Why Alphabetic Keyboards Are Not Easy to Use: Keyboard Layout Doesn't Much Matter*. November 1981.
- 8107. Donald R. Gentner. *Evidence Against a Central Control Model of Timing in Typing*. December 1981.
- 8201. Jonathan T. Grudin and Serge Larochelle. *Digraph Frequency Effects in Skilled Typing*. February 1982.

8202. Jonathan T. Grudin. *Central Control of Timing in Skilled Typing*. February 1982.
8203. Amy Geoffroy and Donald A. Norman. *Ease of Tapping the Fingers in a Sequence Depends on the Mental Encoding*. March 1982.
8204. LNR Research Group. *Studies of Typing from the LNR Research Group: The role of context, differences in skill level, errors, hand movements, and a computer simulation*. May 1982.
8205. Donald A. Norman. *Five Papers on Human-Machine Interaction*. May 1982.
8206. Naomi Miyake. *Constructive Interaction*. June 1982.
8207. Donald R. Gentner. *The Development of Typewriting Skill*. September 1982.
8208. Gary Perlman. *Natural Artificial Languages: Low-Level Processes*. December 1982.
8301. Michael C. Mozer. *Letter Migration in Word Perception*. April 1983.
8302. David E. Rumelhart and Donald A. Norman. *Representation in Memory*. June 1983.
8303. The HMI Project at University of California, San Diego. *User Centered System Design: Papers for the CHI '83 Conference on Human Factors in Computer Systems*. November 1983.
8304. Paul Smolensky. *Harmony Theory: A Mathematical Framework for Stochastic Parallel Processing*. December 1983.

ICS Technical Report List

- 8301. David Zipser. *The Representation of Location*. May 1983.
- 8302. Jeffrey Elman & Jay McClelland. *Speech Perception as a Cognitive Process: The Interactive Activation Model*. April 1983.
- 8303. Ron Williams. *Unit Activation Rules for Cognitive Networks*. November 1983.
- 8304. David Zipser. *The Representation of Maps*. November 1983.
- 8305. The HMI Project. *User Centered System Design: Papers for the CHI '83 Conference on Human Factors in Computer Systems*. November 1983.
- 8306. Paul Smolensky. *Harmony Theory: A Mathematical Framework for Stochastic Parallel Processing*. December 1983.

<p>Navy</p> <p>1 Robert Ahlers Code M711 Human Factors Laboratory NAVTREAE/OU1PCEN Orlando, FL 32813</p> <p>1 Dr. Ed Allen Navy Personnel R&D Center San Diego, CA 92152</p> <p>1 Dr. Mich Bond Office of Naval Research Liaison Office, Far East APO San Francisco, CA 96503</p> <p>1 Dr. Richard Cantone Navy Research Laboratory Code 7510 Washington, DC 20375</p> <p>1 Dr. Fred Chang Navy Personnel R&D Center San Diego, CA 92152</p> <p>1 Dr. Stanley Collier Office of Naval Technology 800 N. Quincy Street Arlington, VA 22217</p> <p>1 CDR Mike Curran Office of Naval Research 800 N. Quincy St. Code 270 Arlington, VA 22217</p> <p>1 Dr. PAT FEDERICO Code P13 NRDC San Diego, CA 92152</p> <p>1 Dr. Jude Franklin Code 7511 Navy Research Laboratory Washington, DC 20375</p> <p>1 Dr. Jim Hollan Code 14 Navy Personnel R & D Center San Diego, CA 92152</p> <p>1 Dr. Ed Hutchins Navy Personnel R&D Center San Diego, CA 92152</p>	<p>Navy</p> <p>1 Dr. Peter Kincaid Training Analysis & Evaluation Group Dept. of the Navy Orlando, FL 32813</p> <p>1 Dr. Joe McLachlan Navy Personnel R&D Center San Diego, CA 92152</p> <p>1 Dr. James McMichael Navy Personnel R&D Center San Diego, CA 92152</p> <p>1 Dr William Montague NRDC Code 13 San Diego, CA 92152</p> <p>1 Technical Director Navy Personnel R&D Center San Diego, CA 92152</p> <p>6 Commanding Officer Naval Research Laboratory Code 2627 Washington, DC 20390</p> <p>1 Office of Naval Research Code 433 800 N. Quincy Street Arlington, VA 22217</p> <p>1 Office of Naval Research Code 44NP 800 N. Quincy Street Arlington, VA 22217</p> <p>6 Personnel & Training Research Group Code 44ZPT Office of Naval Research Arlington, VA 22217</p> <p>1 Psychologist ONR Branch Office 1030 East Green Street Pasadena, CA 91101</p> <p>1 Dr. Bernard Rinaland (OIC) Navy Personnel R&D Center San Diego, CA 92152</p> <p>1 Dr. Robert G. Smith Office of Chief of Naval Operations OP-987H Washington, DC 20350</p>	<p>Navy</p> <p>1 Dr. Alfred P. Smode, Director Training Analysis & Evaluation Group Dept. of the Navy Orlando, FL 32813</p> <p>1 Dr. Richard Snow Liaison Scientist Office of Naval Research Branch Office, London Box 39 PPO New York, NY 09510</p> <p>1 Dr. Richard Sorensen Navy Personnel R&D Center San Diego, CA 92152</p> <p>1 Dr. Frederick Steinhilber Navy Annex CNO - OP115 Arlington, VA 20370</p> <p>1 Roger Weissinger-Baylon Department of Administrative Sciences Naval Postgraduate School Monterey, CA 93940</p> <p>1 Mr John H. Wolfe Navy Personnel R&D Center San Diego, CA 92152</p> <p>1 Dr. Wallace Wulfeck, III Navy Personnel R&D Center San Diego, CA 92152</p>	<p>Marine Corps</p> <p>1 W. William Greenup Education Advisor (E0311) Education Center, MCEBC Quantico, VA 22134</p> <p>1 Special Assistant for Marine Corps Matters Code 100H Office of Naval Research 800 N. Quincy St. Arlington, VA 22217</p> <p>1 DR. A.L. SLAFKOSKY SCIENTIFIC ADVISOR (CODE RD-1) HQ, U.S. MARINE CORPS WASHINGTON, DC 20380</p>
--	--	--	---

Army

- 1 Technical Director
U. S. Army Research Institute for the Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333
- 1 Dr. Beatrice J. Farr
U. S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333
- 1 Dr. Harold F. O'Neil, Jr.
Director, Training Research Lab
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333
- 1 Commander, U.S. Army Research Institute for the Behavioral & Social Sciences
ATTN: PERI-RR (Dr. Judith Orsano)
5001 Eisenhower Avenue
Alexandria, VA 22333
- 1 Joseph Foutla, Ph.D.
ATTN: PERI-1C
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333
- 1 Dr. Robert Seemor
U. S. Army Research Institute for the Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Air Force

- 1 Technical Documents Center
Air Force Human Resources Laboratory
WPAFB, OH 45433
- 1 U.S. Air Force Office of Scientific Research
Life Sciences Directorate, ML
Bolling Air Force Base
Washington, DC 20332
- 1 Air University Library
ARL/LSF 76/443
Maxwell AFB, AL 36112
- 1 Bryan Dallman
AFPHL/LRT
Lowry AFB, CO 80230
- 1 Dr. Alfred R. Fregly
AFOSR/ML
Bolling AFB, DC 20332
- 1 Dr. Genevieve Maddad
Program Manager
Life Sciences Directorate
AFOSR
Bolling AFB, DC 20332
- 1 Dr. T. M. Longridge
AFPHL/GTE
Williams AFB, AZ 85224
- 1 Dr. John Tangney
AFOSR/ML
Bolling AFB, DC 20332
- 1 Dr. Joseph Yasutake
AFPHL/LRT
Lowry AFB, CO 80230

Department of Defense

- 12 Defense Technical Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC
- 1 Military Assistant for Training and Personnel Technology
Office of the Under Secretary of Defense for Research & Engineering
Room 3D129, The Pentagon
Washington, DC 20301
- 1 Major Jack Thorpe
DARPA
1400 Wilson Blvd.
Arlington, VA 22209
- 1 Dr. Robert A. Wisner
OUSDRE (EIS)
The Pentagon, Room 3D129
Washington, DC 20301

Civilian Agencies

- 1 Jr. Susan Chipman
Learning and Development
National Institute of Education
1200 19th Street NW
Washington, DC 20208
- 1 Chief, Psychological Research Branch
U. S. Coast Guard (G-P-1/2/7742)
Washington, DC 20593
- 1 Dr. Edward C. Weiss
National Science Foundation
1800 G Street, NW
Washington, DC 20550
- 1 Dr. Joseph L. Young, Director
Memory & Cognitive Processes
National Science Foundation
Washington, DC 20550

Private Sector

- 1 Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213
- 1 Dr. Alan Beddeley
Medical Research Council
Applied Psychology Unit
15 Chaucer Road
Cambridge CB2 2EF
ENGLAND
- 1 Dr. Patricia Baggett
Department of Psychology
University of Colorado
Boulder, CO 80309
- 1 Mr. Avron Barr
Department of Computer Science
Stanford University
Stanford, CA 94305
- 1 Dr. Menucha Birenbaum
School of Education
Tel Aviv University
Tel Aviv, Ramat Aviv 69978
Israel
- 1 Dr. John Black
Yale University
Box 11A, Yale Station
New Haven, CT 06520
- 1 Dr. John S. Brown
XEROX Palo Alto Research Center
3333 Coyote Road
Palo Alto, CA 94304
- 1 Dr. Glenn Bryan
6208 Poe Road
Bethesda, MD 20817
- 1 Dr. Bruce Buchanan
Department of Computer Science
Stanford University
Stanford, CA 94305
- 1 Dr. Jaime Carbonell
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Private Sector

- 1 Dr. Pat Carpenter
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213
- 1 Dr. William Chase
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213
- 1 Dr. Micheline Chi
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213
- 1 Dr. William Clancey
Department of Computer Science
Stanford University
Stanford, CA 94306
- 1 Dr. Michael Cole
University of California
at San Diego
Laboratory of Comparative
Human Cognition - D003A
La Jolla, CA 92093
- 1 Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138
- 1 Dr. Emanuel Donchin
Department of Psychology
University of Illinois
Champaign, IL 61820
- 1 Dr. Jeffrey Elman
University of California, San Diego
Department of Linguistics
La Jolla, CA 92093
- 1 ERIC Facility-Acquisitions
4833 Rugby Avenue
Bethesda, MD 20014
- 1 Mr. Wallace Feurzeir
Department of Educational Technology
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02238

Private Sector

- 1 Dr. Dexter Fletcher
University of Oregon
Department of Computer Science
Eugene, OR 97403
- 1 Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138
- 1 Dr. Michael Genesereth
Department of Computer Science
Stanford University
Stanford, CA 94305
- 1 Dr. Don Gentner
Center for Human Information Processing
University of California, San Diego
La Jolla, CA 92093
- 1 Dr. Dedre Gentner
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02138
- 1 Dr. Robert Glaser
Learning Research & Development Center
University of Pittsburgh
3939 O'Hara Street
PITTSBURGH, PA 15260
- 1 Dr. Joseph Goguen
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025
- 1 Dr. Daniel Gopher
Faculty of Industrial Engineering
& Management
TECHNION
Haifa 32000
ISRAEL
- 1 DR. JAMES G. GREENO
LRDC
UNIVERSITY OF PITTSBURGH
3939 O'HARA STREET
PITTSBURGH, PA 15213
- 1 Dr. Barbara Hayes-Roth
Department of Computer Science
Stanford University
Stanford, CA 95305

Private Sector

- 1 Dr. Frederick Hayes-Roth
Technolodge
523 University Ave.
Palo Alto, CA 94301
- 1 Dr. Joan I. Miller
Graduate Group in Science and
Mathematics Education
c/o School of Education
University of California
Berkeley, CA 94720
- 1 Dr. James R. Hoffman
Department of Psychology
University of Delaware
Newark, DE 19711
- 1 American Institutes for Research
1055 Thomas Jefferson St., N.W.
Washington, DC 20007
- 1 Dr. Earl Hunt
Dept. of Psychology
University of Washington
Seattle, WA 98105
- 1 Dr. Marcel Just
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213
- 1 Dr. Steven W. Keele
Dept. of Psychology
University of Oregon
Eugene, OR 97403
- 1 Dr. Scott Kelso
Haskins Laboratories, Inc
270 Crown Street
New Haven, CT 06510
- 1 Dr. David Kieras
Department of Psychology
University of Arizona
Tucson, AZ 85721
- 1 Dr. Walter Kintsch
Department of Psychology
University of Colorado
Boulder, CO 80302

Private Sector

- 1 Dr. Stephen Kosllyn
1236 William James Hall
33 Kirkland St.
Cambridge, MA 02138
- 1 Dr. Pat Langley
The Robotics Institute
Carnegie-Mellon University
Pittsburgh, PA 15213
- 1 Dr. Jill Larkin
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213
- 1 Dr. Alan Lesgold
Learning M&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260
- 1 Dr. Jim Levin
University of California
at San Diego
Laboratory for Comparative
Human Cognition - 0003A
La Jolla, CA 92093
- 1 Dr. Michael Levine
Department of Educational Psychology
210 Education Bldg.
University of Illinois
Champaign, IL 61801
- 1 Dr. Don Lyon
AFRL/OT (DDR1)
Williams AFB, AZ 85225
- 1 Dr. Jay McClelland
Department of Psychology
MIT
Cambridge, MA 02139
- 1 Dr. James B. Miller
ComputerThought Corporation
1721 West Plano Parkway
Plano, TX 75775
- 1 Dr. Mark Miller
ComputerThought Corporation
1721 West Plano Parkway
Plano, TX 75775

Private Sector

- 1 Dr. Allen Muro
Behavioral Technology Laboratories
1845 Elena Ave., Fourth Floor
Redondo Beach, CA 90277
- 1 Dr. Donald A. Norman
Cognitive Science, C-015
Univ. of California, San Diego
La Jolla, CA 92093
- 1 Dr. Jesse Orlinsky
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311
- 1 Dr. Nancy Pennington
University of Chicago
Graduate School of Business
1101 E. 58th St.
Chicago, IL 60637
- 1 DR. PETER POLSON
DEPT. OF PSYCHOLOGY
UNIVERSITY OF COLORADO
BOULDER, CO 80509
- 1 Dr. Mike Posner
Department of Psychology
University of Oregon
Eugene, OR 97403
- 1 Dr. Fred Ruff
Physics Department
University of California
Berkeley, CA 94720
- 1 Dr. Lauren Bessick
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 1521
- 1 Mary S. Riley
Program in Cognitive Science
Center for Human Information Processing
University of California, San Diego
La Jolla, CA 92093
- 1 Dr. Andrew M. Rose
American Institutes for Research
1055 Thomas Jefferson St. NW
Washington, DC 20007

Private Sector

- 1 Dr. Ernst Z. Rothkopf
Bell Laboratories
Murray Hill, NJ 07974
- 1 Dr. William B. Rouse
Georgia Institute of Technology
School of Industrial & Systems
Engineering
Atlanta, GA 30332
- 1 Dr. David Rumlhart
Center for Human Information Processing
Univ. of California, San Diego
La Jolla, CA 92093
- 1 Dr. Michael J. Smet
Perceptronics, Inc
6271 Varial Avenue
Woodland Hills, CA 91364
- 1 Dr. Arthur Samuel
Yale University
Department of Psychology
Box 11A, Yale Station
New Haven, CT 06520
- 1 Dr. Roger Schank
Yale University
Department of Computer Science
P.O. Box 2158
New Haven, CT 06520
- 1 Dr. Walter Schneider
Psychology Department
603 E. Daniel
Champaign, IL 61820
- 1 Mr. Colin Sheppard
Applied Psychology Unit
Admiralty Marine Technology Est.
Teddington, Middlesex
United Kingdom
- 1 Dr. M. Wallace Simaiko
Program Director
Hempner Research and Advisory Services
Switcheon Institution
801 North Pitt Street
Alexandria, VA 22314
- 1 Dr. Edward E. Smith
Bell Beranek & Newman, Inc.
50 Hamilton Street
Cambridge, MA 02138

Private Sector

- 1 Dr. Elliott Soloway
Yale University
Department of Computer Science
P.O. Box 2158
New Haven, CT 06520
- 1 Dr. Kathryn T. Spohr
Psychology Department
Brown University
Providence, RI 02912
- 1 Dr. Robert Sternberg
Dept. of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520
- 1 Dr. Albert Stevens
Bell Beranek & Newman, Inc.
10 Hamilton St.
Cambridge, MA 02238
- 1 David E. Stone, Ph.D.
Resaltine Corporation
7400 Old Springhouse Road
McLean, VA 22102
- 1 Dr. Eihumi Tateoka
Computer Based Education Research Lab
252 Engineering Research Laboratory
Erbana, IL 61801
- 1 Dr. Maurice Tateoka
230 Education Bldg
1310 S. Sixth St.
Champaign, IL 61820
- 1 Dr. Perry V. Thorndyke
Perceptronics, Inc.
545 Middlefield Road, Suite 140
Menlo Park, CA 94025
- 1 Dr. Douglas Touss
Univ. of So. California
Behavioral Technology Lab
1845 S. Elena Ave.
Redondo Beach, CA 90277
- 1 Dr. Earl Van Lahn
Zerax P&C
3333 Coyote Hill Road
Palo Alto, CA 94304

Private Sector

1 Dr. Keith T. Weccourt
Perceptics, Inc.
545 Middlefield Road, Suite 140
Menlo Park, CA 94025

1 William B. Whitton
Bell Laboratories
2B-610
Holmdel, NJ 07733

1 Dr. Christopher Wichens
Department of Psychology
University of Illinois
Champaign, IL 61820

1 Dr. Thomas Wichens
Department of Psychology
Frama Hall
University of California
405 Hilgard Avenue
Los Angeles, CA 90024

1 Dr. Mike Williams
IntelliGenetics
124 University Avenue
Palo Alto, CA 94301

1 Dr. Joseph Wohl
Alphatech, Inc.
2 Burlington Executive Center
111 Middlesex Turnpike
Burlington, MA 01803

END

FILMED

3-84

DTIC