

NO-A104 039

CATEGORY ACCESSIBILITY EFFECTS IN A SIMULATED  
EXEMPLAR-BASED MEMORY(U) PURDUE UNIV LAFAYETTE IN DEPT  
OF PSYCHOLOGICAL SCIENCES E R SMITH 1987 TR-ONR-5  
N00014-84-K-0200


1/1

UNCLASSIFIED

F/G 5/0

NL





SECURITY

TATION PAGE

Form Approved  
OMB No. 0704-0188

AD-A184 839

1a REPC

1b RESTRICTIVE MARKINGS

2a SECURITY CLASSIFICATION AUTHORITY

3 DISTRIBUTION / AVAILABILITY OF REPORT

Approved for public release;  
distribution unlimited

2b DECLASSIFICATION / DOWNGRADING SCHEDULE

4. PERFORMING ORGANIZATION REPORT NUMBER(S)

TR-ONR-5

5 MONITORING ORGANIZATION REPORT NUMBER(S)

6a. NAME OF PERFORMING ORGANIZATION  
Purdue University

6b OFFICE SYMBOL  
(If applicable)

OXD

7a. NAME OF MONITORING ORGANIZATION

Cognitive Science Program  
Office of Naval Research (Code 1142CS)

6c. ADDRESS (City, State, and ZIP Code)

Dept. of Psychological Sciences  
W. Lafayette, IN 47907

7b. ADDRESS (City, State, and ZIP Code)

800 North Quincy St.  
Arlington, VA 22217-5000

8a. NAME OF FUNDING / SPONSORING ORGANIZATION

8b OFFICE SYMBOL  
(If applicable)

9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER

N00014-84-K-0288

8c. ADDRESS (City, State, and ZIP Code)

10 SOURCE OF FUNDING NUMBERS

PROGRAM ELEMENT NO	PROJECT NO	TASK NO.	WORK UNIT ACCESSION NO
61153N			

11 TITLE (Include Security Classification)

Category accessibility effects in a simulated exemplar-based memory

12 PERSONAL AUTHOR(S)

Smith, Eliot R.

13a. TYPE OF REPORT

Interim technical

13b TIME COVERED

FROM \_\_\_\_\_ TO \_\_\_\_\_

14 DATE OF REPORT (Year, Month, Day)

April 1987

15 PAGE COUNT

16

16 SUPPLEMENTARY NOTATION

In press, Journal of Experimental Social Psychology

17 COSATI CODES

FIELD	GROUP	SUB-GROUP
05	10	

18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)

memory, social cognition, social categorization

19 ABSTRACT (Continue on reverse if necessary and identify by block number)

A quantitative model of long-term memory is applied, in the form of a computer simulation, in an attempt to reproduce several known properties of social priming or category accessibility effects. This exemplar-only model, developed by Hintzman, involves the storage and retrieval of specific experiences, without any representation of abstract schemas, constructs, or prototypes. It successfully reproduces several basic findings concerning effects of priming on category accessibility, including the cumulative impact of several primes, the greater influence of priming with behavior descriptions versus trait labels, and the over-time change in the relative impact of frequent versus recent priming. The simulation not only accurately reproduces current findings, but also makes new predictions that can be empirically tested. The simulation is useful in evaluating the extent to which a simple model of the storage of specific experiences can account for these effects. Conversely, it may also permit

OVER

20 DISTRIBUTION / AVAILABILITY OF ABSTRACT

UNCLASSIFIED/UNLIMITED  SAME AS RPT  DTIC USERS

21 ABSTRACT SECURITY CLASSIFICATION

unclassified

22a NAME OF RESPONSIBLE INDIVIDUAL

Dr. James Lester

22b TELEPHONE (Include Area Code)

202-696-4503

22c OFFICE SYMBOL

ONR 1142CS

19. identification of points at which such an account is inadequate and requires supplementation by other processes.

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	20



Category Accessibility Effects in a Simulated Exemplar-Based Memory

ELIOT R. SMITH  
Purdue University

Received July 14, 1986

A quantitative model of long-term memory is applied, in the form of a computer simulation, in an attempt to reproduce several known properties of social priming or category accessibility effects. This exemplar-only model, developed by Hinzman, involves the storage and retrieval of specific experiences, without any representation of abstract schemas, constructs, or prototypes. It successfully reproduces several basic findings concerning effects of priming on category accessibility, including the cumulative impact of several primes, the greater influence of priming with behavior descriptions versus trait labels, and the over-time change in the relative impact of frequent versus recent priming. The simulation not only accurately reproduces current findings, but also makes new predictions that can be empirically tested. The simulation is useful in evaluating the extent to which a simple model of the storage of specific experiences can account for these effects. Conversely, it may also permit identification of points at which such an account is inadequate and requires supplementation by other processes. © 1987 Academic Press, Inc.

Category accessibility or social priming effects (Higgins, Rholes, & Jones, 1977; Srull & Wyer, 1979) have profoundly influenced theoretical development in social cognition, perhaps most evidently in the well-known Wyer and Srull model (1980, 1986) and in Higgins' work on accessibility of constructs in memory (Higgins & King, 1981).

Category accessibility effects have been demonstrated by research in two basic paradigms. Srull and Wyer (1979) use "scrambled sentences" describing behaviors to prime a construct such as hostility. Under a cover story concerning "the way people perceive word relationships," subjects read sets of words like *leg break arm his* and underline three words that make a complete sentence. In the relevant experimental conditions, a large proportion of the sentences that can be created have

This research was supported by the Office of Naval Research under Contract N00014-84-K-0288. Correspondence regarding this article (including requests for reprints) should be sent to Eliot R. Smith, Department of Psychological Sciences, Purdue University, West Lafayette IN 47907, or to ESMITH@PURCCVM via BITNET.

0022-1011/87 \$3.00

Copyright © 1987 by Academic Press, Inc.

All rights of reproduction in any form reserved.

0022-1011/87 \$3.00

Copyright © 1987 by Academic Press, Inc.

All rights of reproduction in any form reserved.

—DESIGN 1AP/JAN 1987—BUJ PACMS 310855224 06-26-87 11-12-88—

4  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
352  
353  
354  
355  
356  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431  
432  
433  
434  
435  
436  
437  
438  
439  
440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458  
459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482  
483  
484  
485  
486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500

87 9 14 017

content related to hostility, as in the example. Next, believing that they are participating in an unrelated experiment, subjects read a paragraph describing a fictitious character's behaviors, which are ambiguously hostile. They rate "Donald" on several scales related to hostility, which constitute the dependent measure. The essential finding from these studies is that priming with hostility-related materials increases subjects' ratings of the target character's hostility. The effect can be obtained as much as 24 h after the priming manipulation (Sruil & Wyer, 1979). The theoretical explanation for the effect is that priming puts the construct of hostility at or near the top of its Storage Bin in memory (Wyer & Sruil, 1980, 1986). When the subjects read about Donald's behaviors, the hostile construct is therefore more accessible and more likely to be used in interpreting them because Storage Bins are searched from the top down until an applicable construct is located.

Higgins and his associates have used a related paradigm (e.g., Higgins et al., 1977; Higgins, Bargh, & Lombardi, 1985). Under a cover story these researchers expose subjects to trait words related to the target construct (e.g., reckless or adventuresome). The subjects then read descriptions of a target character's behaviors that are ambiguously related to the primed constructs and rate the target character. Category accessibility is influenced by priming in this paradigm also, though the effect has only been shown to last a matter of minutes, not 24 h. Higgins et al. (1985) have developed a theoretical model in which constructs in memory are associated with an energy cell, whose charge increases with priming but decreases with the passage of time. The applicable construct associated with the highest charge is selected for use in interpreting behaviors.

Though the Wyer and Sruil and Higgins models differ in many respects, their similarities are more significant for the purposes of this paper. The models agree that (a) abstract representations of constructs are stored in memory (perhaps in the form of schemas or prototypes), and (b) priming influences some property of those representations (e.g., position in a Storage Bin, or charge of an energy cell). This paper uses a computer simulation method to evaluate a theoretical alternative to both of these assumptions. I will show that many properties of category accessibility effects can be accounted for using a simple model in which (a) no constructs or other abstract representations—only records of specific episodes—are stored in memory, and (b) the effect of priming is only to add new episodic records to memory, rather than to change the accessibility or other properties of existing constructs. The model is essentially that of Hintzman (1984, 1986).

Hintzman's model has successfully simulated a number of findings in the area of memory and categorization. Most relevant for the concerns of this paper is its success with findings that have traditionally been taken as evidence for the abstraction and storage of category prototypes in

memory. For example, after learning a category by studying a series of category names paired with example category members (which have varying degrees of similarity to a prototype), people can categorize the category prototype though they have never seen it before. Moreover, people can categorize the prototype faster and more confidently than they can old, familiar category examples, and the ability to categorize the prototype may be more stable over time (i.e., more resistant to forgetting). Such findings had been taken as evidence for a memory model in which, along with some record of encounters with individual category members, a representation of the typical or average category member is abstracted and stored as a schema or prototype. Hintzman's (1984, 1986) demonstrations, in which an exemplar-only memory reproduces these effects, suggest the possibility of applying the same model to category accessibility effects, which also have traditionally been viewed as evidence that memory includes abstract constructs.

The Hintzman Memory Model

Hintzman's (1986) model is quite simple in its basic structure. The perceiver is assumed to be sensitive to a large number of properties (features) of experiences, some of which may be closely linked to sensory inputs while others may be abstract. Examples might be features representing the color, spatial location, or size of an object. Each episode or experience is encoded and represented in memory by a number of binary features, which take the values 1 or -1. The value 0 is also possible, and means that the value of a particular feature is indeterminate. Different subsets of features within the memory trace may be assigned to represent different aspects of the experience, such as the particular stimulus object encountered, its category label, and the context in which it was encountered.

A basic assumption of the memory model is that a record of every experience is stored in memory, even if it is highly similar to previous experiences. Storing an experience amounts to copying the set of features into an array in which all long-term memory (LTM) traces are preserved (not necessarily perfectly). Each memory trace therefore records a single episode or experience. Forgetting under the model is a process in which randomly chosen features in the LTM array are changed to zero.

Retrieval is treated as a process in which a response is elicited from memory by a probe, which can be either a complete array of features or an array with a subset of features set to zero (indicating unknown). The former case describes recognition memory; the memory response can be used to judge whether or not an experience similar to the probe has already been stored. The latter case describes associative learning, for the information retrieved from memory may include values of the

features that were unspecified in the probe. This aspect of the model is used here to simulate learning and retrieval of behavior-trait relationships. The model's generation of a response to a probe is described by a few equations (see Appendix A). An intuitive interpretation of the process is as follows. All stored traces are activated by a probe, to an extent depending on each trace's similarity to the probe. Similarity is a function of the proportion of feature values that the trace shares with the probe. Each trace in LTM contributes to the overall memory response, which is the summation of the content of every individual trace, weighted by the trace's activation level. These assumptions mean that traces that are similar to the probe will be more highly activated and will contribute most heavily to the memory response. If those traces include values for features where the probe had zeros, the memory response will include that information.

One of Hintzman's simulations will illustrate the model's operation. This example applies the model to a category-learning or "schema-abstraction" task (Hintzman, 1986), which is related to the category-accessibility studies to be considered below. In such a task (e.g., Homa, Cross, Cornell, Goldman, & Schwartz, 1973), subjects are trained to classify patterns (which could be dot patterns, letter strings, or multiauttribute descriptions of objects) into several categories. The patterns in each category are constructed to be similar to a single category prototype which subjects do not actually see during training. After training, subjects are tested with other patterns, and are required to assign each to a category. Test patterns can include the category prototypes, the old patterns used in the training, new patterns that are similar to the prototypes, or completely random patterns.

To simulate this task, Hintzman (1986) randomly generates a series of features representing a prototype pattern and a category label for each category to be used. Several distortions of each prototype are then formed by altering randomly selected features, to generate a stimulus with some degree of similarity to the prototype. Distortions are paired with the category label and stored in LTM. This corresponds to the training phase. For testing, a test pattern is constructed (as a copy of the prototype, one of the training patterns, a new distortion, etc.) and presented as a memory probe, with the features corresponding to its category label given zero (unknown) values. This corresponds to the subject's classification task presented with an unlabeled exemplar; they have to assign it to a category. The response from the memory will include nonzero values in the feature positions corresponding to the name. The basic datum is the similarity (expressed as a correlation across the label features) between the model's output and the actual label of the correct category, or the probability of assigning the test probe to the correct category. Hintzman, in a series of simulations that follow this general outline, shows that

174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215  
216  
11

several basic facts about human performance in this category-learning paradigm are qualitatively reproduced by his model.

The simulations of category-accessibility studies to be reported in this paper follow the same general outline. Here, the category is thought of as a trait (e.g., hostility) to which a particular behavior (an instance of the category) is to be assigned. An initial training phase represents preexperimental associations between particular types of behaviors and trait categories, by storing several behavior-trait pairs from each of several categories. Priming then takes place: new traces related to the to-be-primed category are stored in memory. Finally the subject is tested with an ambiguous behavior—one that is ambiguously related to the category—and the model's response (the trait category to which it is assigned) is recorded. The procedure for the simulation is described in more detail below, following the presentation of some hypotheses.

*Evidence Regarding Category Accessibility*

Three empirical generalizations regarding category accessibility are briefly described here. In this paper, the goal is to see whether these particular data patterns can be reproduced by the simulation, in a qualitative sense. As Hintzman notes, "It would be inappropriate to assume that the model somehow captures all of the variables that contribute significantly to the performance of human subjects in an experiment, and so no attempt has been made to fit data quantitatively. Rather, the simulations were intended to investigate the functional relationships that the model predicts under the manipulation of a variety of experimental variables." (1986, p. 413)

1. *Number of primes.* The most basic finding is that the number of primes (i.e., frequency of encountering category-related information) influences category accessibility. Srull and Wyer (1979) and many other studies have replicated this effect. The prediction is that increasing the number of primes will increase the probability of classifying the test stimuli using the primed category.

2. *Category label primes versus behavior primes.* Smith and Branscombe (1986) exposed subjects to equal numbers of primes which were either trait-related words (like *hostile*, *unfriendly*, in effect, instances of the name of the to-be-primed category as used in the Higgins priming paradigm) or behaviors (the Srull & Wyer hostility-related word sets). I assume (as do Wyer & Srull, 1980) that subjects presented with behaviors access the relevant trait, perhaps without being consciously aware of it. Smith and Miller (1983) and Winter and Uleman (1984) have provided evidence for such a process of spontaneous trait inference from behaviors. Therefore the record stored in memory is not simply the behavioral instance, but the instance plus the category label. The Smith and Branscombe results show that priming with instances of behaviors has longer lasting effects




Fig. 1. Illustration of different types of memory traces. Note: Underlined features differ from the prototype.

253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
300  
301  
302  
11

ELIOT R. SMITH

Category	Name	Instance	Context
Standard exemplar	+	+	+
Label primes	+	+	+
Behavior primes	+	+	+
Primes	+	+	+

than priming with trait-related words. Other research converges with this result, for example, Snull and Wyer (1979) find priming over a 24-h period, though studies in the Higgins paradigm generally use a delay of only a few minutes between priming and test.

3. *Decay of effects of varying numbers of primes.* Higgins et al. (1985) demonstrated that priming effects decay more slowly over time if the category was primed more often. Their data showed neither pure recency effects (the most recently primed constant remaining the most active) nor pure frequency effects (the most frequently primed construct dominating) but a mixture of the two. Specifically, they consider the situation where one category is primed several times and then after some delay another category is primed just once. On an immediate test, the recently primed category may be the more active. However, as time passes the more recently primed category will have activation more rapidly than the other, more frequently primed category will predominate.

This simulation then examines the effects of varying the number of primes, type of primes (category labels versus behaviors), and the delay between priming and test (operationalized as the amount of forgetting that intervenes). The results are compared to the relevant findings from the above experiments with human subjects, to see how well the simple exemplar-only memory model can reproduce them.

METHOD

The Hintzman model was described in general terms above. Specific equations describing its operation and a simple example are presented in Appendix A. The application of the model to the category accessibility paradigm proceeded as follows: most of the approach and specific procedure follow Hintzman's (1986) application of the model to the category-learning task.

1. The prime and each memory trace were defined by 27 features. Of these, 10 represent the category name (trait), 11 the specific category instance (behavior), and 4 the context in which the episode was experienced (see Fig. 1).
2. Three categories (traits) were used. Their names and prototype instances (behaviors) were generated at random with values of 1 and -1 equally likely for each feature. The top panel of Fig. 1 shows an example.
3. Example behaviors for training were generated by altering (multiplying by -1) the randomly chosen features of the prototype behavior. Four such distortions of each prototype

were stored in memory with the category name, and with the context features set to 0000. To represent preexperimental knowledge about the traits and their behavioral instances, the second panel of Fig. 1 shows an example, with the altered features underlined. A total of 12 traces were therefore stored in memory in this (training) phase of the simulation.

4. Primes were formed in two ways (see Fig. 1). Label primes were undistorted copies of the category label (the first 10 features), with the instance features set to zero and the context features set to 1111 (representing the experimental context). Behavior primes were generated by randomly distorting two features of the prototype behavior of the target category and storing it with the category name and the 1111 context. This follows our assumption mentioned earlier, that a behavior sentence generally elicits a trait inference from the perceiver. A behavior prime was thus formed in the same way as a training exemplar except for the value of the context features. Depending on the condition being simulated, varying numbers of primes were formed and stored in memory.

5. Depending on the condition, forgetting might take place. The model does not directly represent the passage of time, but assumes that the impact of time on memory performance is mediated by forgetting. Hintzman does not specify the form of the relationship between time and forgetting. I borrow an assumption from Anderson (1983), that memory decays as a power function of time. Anderson shows that this particular functional form fits well with many experimental findings, but for our purposes the exact function is probably unimportant. The simulations in this paper assume that at time  $t$  (greater than or equal to 1), a proportion  $t^{-1}$  of features in LTM remain unchanged. For example, at  $t = 5$  (in arbitrary units), 45% of features remain, implying that 55% of them randomly selected have been replaced with zeros. At  $t = 10$ , 32% of features remain (68% have been "forgotten"), and so on.

6. Final access probe was constructed by changing 4 of the 13 features of the prototype pattern from the target category, as shown in the bottom panel of Fig. 1. As Hintzman notes, this may not appear to be much of a distortion of the prototype, but changing 6.5% of the 13 features would make a pattern that is orthogonal to the prototype and thus completely unrelatable to the category. Four features therefore seem adequate to represent the theoretical notion of an instance that is ambiguously related to the category. The test probe was presented to the memory model with the label features set to zero and the context features set to 1111; testing like priming takes place in the experimental context.

The model's response, the label features in the memory output, was correlated with the three actual category labels. The results of the trial are the correlation across the 10 label features with the target category label and the correct or incorrect category assignment based on the output. The assignment is counted as correct if the correlation with the target category label is positive and greater than the correlation with any other category.

This procedure simulates one subject. The entire procedure from the random generation of category names onward, was repeated for the next subject. A total of 60 experimental conditions was defined by the number of primes (0, 1, 2, 4, 8, 16), type of primes (label, behavior), and postprime forgetting (0, 5, 10, 20, or 40 time units, corresponding to 00, 55, 60, 70, or 84% of features forgotten). Five hundred subjects were simulated in each condition to produce relatively stable results.

RESULTS

The results are presented in Fig. 2 (label primes) and Fig. 3 (behavior primes). The graphed dependent variable is the proportion categorized in the primed category (across the 500 "subjects"), but results based on the average correlation of the memory response with the target category label are qualitatively similar. The standard errors of the graphed points average about .02 (an exact standard error for each point can be computed

*Recent versus Frequent Priming*

The results also show that the decrease in priming effectiveness due to forgetting is slower with a greater number of primes. This pattern was obtained with human subjects by Higgins et al. (1985). For example, the model results show that one behavior prime with no forgetting has a greater effect than two primes after 5 units of time (.866 correctly categorized > .780), but after five further units, the effect of one prime has decreased more and is smaller than two primes (.686 < .720).

**DISCUSSION**

This application of the Hintzman memory model assumes that the effect of priming is simply to add episodic traces to memory, rather than to activate a preexisting abstract construct or schema. In fact, the model does not incorporate abstractions at all. Yet it is able to reproduce the major features of social priming or category accessibility effects. In the model, primes have their effects by mechanisms that are different for behavior and label primes. Each trace representing a behavior-trait pair provides an additional associative link between the behavior and the label, so that when an ambiguous behavior probe is presented without a label, the model is better able to respond with the label. A label prime has no effect because of its link to the experimental context. Other simulations not reported in detail here show that if the context features are always set at 0000, then label priming has no effect.

One question may be raised about this theoretical account. It has been shown that increased category accessibility due to priming is uncorrelated with the perceiver's ability to recall the priming information (Higgins et al., 1985; Smith & Branscombe, 1986). This might be taken as evidence that episodic traces do not underlie category accessibility, for one would expect the traces to be retrievable in the recall task as well, leading to a positive correlation. However, this logic is not valid. The same traces might mediate both recall and category accessibility without the two dependent measures being correlated, because the retrieval cues and retrieval processes required by the two tasks are quite different. Such retrieval differences can lead to independence even between two dependent measures that are known to rely on the same memory structure (Jacoby, 1983; Roediger & Blaxton, in press). For example, recall and recognition memory rest on the same episodic traces yet can be independent (Tulving & Wiseman, 1975). This result is correctly reproduced by Hintzman's model (Hintzman, in press).

The model's ability to reproduce major features of category accessibility effects has implications both for the Hintzman model itself and for theory in social cognition in general. These implications are addressed below, following a brief discussion of the effects of variations in the model and

314  
315  
316  
317  
318  
319  
320  
321  
322



323




Fig. 2. Label primes. Effects of priming and forgetting on classification of ambiguous test probe

by the reader from the binomial formula, if desired). The results are presented in terms of the three specific issues raised above

*Number of Primes*

Clearly, there is an overall priming effect (one or more primes increase category accessibility compared to unprimed subjects) that is greater as more primes are used. This replicates the data pattern of Strull and Waver (1979).

*Behavior versus Trait Priming*

Comparing Fig. 2 with Fig. 3 shows that priming with behavior-trait pairs is more effective, and more resistant to forgetting, than priming with trait labels alone. This replicates Smith and Branscombe (1986).

of some new and testable predictions regarding category accessibility effects that can be derived from this exemplar-only perspective

Model Variations

The simulations reported here incorporate a number of parameters and assumptions, both in the Hintzman model itself and in the way the category accessibility paradigm is represented for the simulation. Different assumptions and different values for the parameters (e.g., the time course of forgetting, the number of features in a memory trace, the method of construction of training, priming, and test stimuli) would produce quantitatively different results, of course. However, for the purpose of this paper, these differences are not crucial. The goal here is to provide an existence proof, a demonstration that an exemplar-only memory model can reproduce certain effects. These results establish that one model from that class can do so, and the fact that other models might or might not do so is irrelevant to this goal. However, it may be worthwhile to discuss the effects of some possible variations in the model. I consider three:

a. All training and priming stimuli in these simulations included undistorted category name features. However, synonyms or alternative labels for the trait category may be encountered in reality, either in preexperimental experience or in priming manipulations. Varying the label features (i.e., using distortions of the prototype label) would weaken priming effects somewhat by introducing random variance into the label features retrieved from the memory. However, the process of forgetting (setting randomly chosen features to zero) also introduces variance into the memory traces, including their label features, and its effects probably differ little from the effects of storing random, distorted features in the first place.

b. The number of features was fixed, following Hintzman's usage (1986), at 10 for the category name, 13 for the instance, and 4 for context. These numbers influence the relative weight of these three parts of a memory trace in the similarity computations. Obviously, two traces that share an instance but differ in context will be more similar than two traces that differ in the instance represented but share context. The effects of varying the numbers of features are predictable from this analysis. For example, increasing the number of context features would make priming effects more context-bound. Priming would have less effect outside of the context in which the priming took place, because the memory traces of the priming events would be less similar to, and less activated by, a probe presented in a different context.

c. These simulations assume that the presentation of a behavior results in the inference of a trait (and therefore in the storage of a behavior-trait pair in LTM), but that the presentation of a trait label does not

result in the inference of a particular behavioral exemplar. It is possible that encountering a trait term (such as hostile) sometimes causes people to think of a corresponding behavior, in which case a behavior-trait pair would be stored in memory. To the extent that this happens, label primes would act more like behavior primes. Since the evidence is that behavior primes have much stronger effects than label primes (Smith & Braascombe, 1986), it seems that the probability of inferring a behavioral instance from a trait label is considerably lower than the probability of making the reverse inference (cf. Smith & Miller, 1983; Winter & Uleman, 1984).

An implication of this line of reasoning is that there are two possible pathways by which a label prime might influence category accessibility. One is via the storage of a behavior-trait pair in memory (an effect that would be the same as that of a behavior prime), and the other is via a link between the label and the context. These simulations set the probability of the first pathway at zero (label primes never elicit corresponding behaviors), in order to show more clearly the effects of context alone in the label-priming conditions. The effects of raising the probability from zero would be to shift the label-prime results toward the behavior-prime results graphed in this paper. If the probability were one, label and behavior primes would have identical effects.

New Predictions

The role of LTM traces of specific experiences in category accessibility effects that is suggested by these results could be further examined by testing two new predictions derived from the Hintzman model. My current claim is only that exemplar-based and schema- or prototype-based models of memory offer equally effective alternative explanations for category accessibility phenomena. Hence, these phenomena do not (as has generally been assumed) provide direct evidence for the existence of schemas or prototypes. Should new predictions like the following be verified, however, the situation would change and a burden would fall on advocates of schema theories to account for the new findings.

a. In the model, the similarity of a new instance (the ambiguous test stimulus) to old, stored instances (the primes) is crucial in determining its categorization. Priming should therefore have the greatest influence on the categorization of test stimuli that are highly similar to the primes, independently of their similarity to the category prototype. Under alternative theoretical viewpoints that stress activation of abstract category prototypes, only the similarity of the test stimulus to the prototype should determine the magnitude of priming effects. Since priming is said to operate by activating a representation of the category as a whole, not by storing specific representations of the priming stimuli. Put another way, the issue is whether the effect of a prime is relatively specific to stimuli that resemble the prime, or whether it is more general, increasing

59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

96

97


98

99


100

101

11



146




an experimental context is followed by dependent measures in a different context. What constitutes the effective context for human subjects is unclear; the physical location and identity of the experimenter are probably less important than the subject's awareness of being in an experiment. To approximate an unrelated-context test, a priming manipulation could be administered in the lab and then subjects called at home later under the pretext of a survey, in which the dependent measure questions are embedded. It will probably be difficult to construct an unrelated-context test of the effects of label priming, because the effects have only been shown to last a matter of minutes even without a context change (Higgins et al., 1985).

Implications

Does this simulation demonstrate that Hintzman's memory model is correct? Though the results count as support for the model, which was not developed with priming or category accessibility effects in mind, that is not the real purpose of this paper. Hintzman's model has many omissions, including procedures for encoding, inference, and the like, which can also change with experience, independent of the contents of declarative memory (Smith & Branscombe, 1996). The model may well prove adequate or be falsified by experimental data gathered in different paradigms. Regardless of its eventual fate, however, the simulations reported here establish one conclusion that is *not* subject to empirical falsification: That is that an exemplar-only memory model can reproduce major features of social category-accessibility effects. The results reported here demonstrate that one member of this class of models has this ability. The point of this conclusion is that category accessibility effects can no longer be taken as evidence for the use of abstract schemas or constructs in social perception. Logically, this resembles Hintzman's demonstration that the exemplar-only memory model can reproduce effects that have been taken as evidence for prototypes in category learning. It may still be true that we have schemas or abstract constructs—in fact, I suspect that we do—but category accessibility effects do not prove it. As Hintzman notes, one benefit of this simulation will be to force more theoretical specificity concerning the role of schemas in social cognition.

Even a theorist who believes that abstract representations play a central role in category accessibility effects must be concerned to know which phenomena require them for their explanation. This may best be determined by a theoretical exercise in which one attempts to get along without abstract traces. From the residual phenomena not explained by the exercise, one should get a hint as to what kind of abstraction process needs to be added. A conclusion of the present study is that no clear examples of such residual phenomena have been uncovered so far, even among very likely candidates, and so their number may be surprisingly small (1986, pp. 422-423).

147



the accessibility of the primed category with respect to a wide range of test stimuli.

The simulation was used to verify that this prediction does follow from the assumptions made in this paper. Four behavior primes were simulated followed by forgetting of .55 of the features (corresponding to 5 time units) in three conditions, with 500 subjects each. Ceiling effects were observed in these runs if just two features were distorted in the primes, so four distortions were used in the primes as well as in the test stimulus. Three conditions were run. With no primes, the model produced .596 correct categorization. With independent randomly chosen primes and probe, each with four features distorted, the proportion correct was .862. However, when the probe was identical to one of the primes, the proportion was .952, significantly higher. Thus, in the model, similarity to a prime strongly increases the priming effect even when similarity to the prototype is held constant (the probe's differed from the prototype in four features in each of these conditions).

Whittlesea (1987) tested hypotheses like these in studies using perceptual fluency dependent variable. Prior exposure to letter-string stimuli drawn from particular categories increased his subjects' ability to perceive similar stimuli in brief (30-ms), masked visual presentations. Whittlesea found that similarity of the test stimuli to previous, encountered instances, rather than their similarity to the category prototype, determined perceptual performance. Similar hypotheses seem not to have been tested for category accessibility dependent variables, but they can and should be.

In the model, the effect of priming depends on a common context between priming and test, to a greater extent for label primes than for behavior primes. This is because for label primes, the label-context link is the only way in which a prime helps the later presentation of a behavior in the same context to activate and hence retrieve the appropriate label. Again, the simulation was used to verify that this prediction about context sensitivity actually follows from the assumptions made here. Four primes followed by .78 forgetting (20 time units) were used. Preexperimental training used a .0000 (unspecified) context and priming a .1111 (experimental) context, as in the main simulations. The context of the test probe varied. With the same .1111 context for the probe, proportion correct was .560 for label primes and .698 for behavior primes (compared to .460 for a no-prime baseline condition). With an unspecified .0000 context for test, results were .426 for label primes and .586 for behavior primes. Finally, with an orthogonal 1, -1, 1, -1 context for test, results were .406 and .572. These results clearly show that removal from the experimental context (to either a general, context-unspecified situation or an unrelated situation) destroys any effects of label priming. Change in context diminishes but does not destroy the effects of behavior priming.

This hypothesis could be tested by experiments in which priming in

148




149



Even a theorist who believes that abstract representations play a central role in category accessibility effects must be concerned to know which phenomena require them for their explanation. This may best be determined by a theoretical exercise in which one attempts to get along without abstract traces. From the residual phenomena not explained by the exercise, one should get a hint as to what kind of abstraction process needs to be added. A conclusion of the present study is that no clear examples of such residual phenomena have been uncovered so far, even among very likely candidates, and so their number may be surprisingly small (1986, pp. 422-423).

150




Does this simulation demonstrate that Hintzman's memory model is correct? Though the results count as support for the model, which was not developed with priming or category accessibility effects in mind, that is not the real purpose of this paper. Hintzman's model has many omissions, including procedures for encoding, inference, and the like, which can also change with experience, independent of the contents of declarative memory (Smith & Branscombe, 1996). The model may well prove adequate or be falsified by experimental data gathered in different paradigms. Regardless of its eventual fate, however, the simulations reported here establish one conclusion that is *not* subject to empirical falsification: That is that an exemplar-only memory model can reproduce major features of social category-accessibility effects. The results reported here demonstrate that one member of this class of models has this ability. The point of this conclusion is that category accessibility effects can no longer be taken as evidence for the use of abstract schemas or constructs in social perception. Logically, this resembles Hintzman's demonstration that the exemplar-only memory model can reproduce effects that have been taken as evidence for prototypes in category learning. It may still be true that we have schemas or abstract constructs—in fact, I suspect that we do—but category accessibility effects do not prove it. As Hintzman notes, one benefit of this simulation will be to force more theoretical specificity concerning the role of schemas in social cognition.

151



Even a theorist who believes that abstract representations play a central role in category accessibility effects must be concerned to know which phenomena require them for their explanation. This may best be determined by a theoretical exercise in which one attempts to get along without abstract traces. From the residual phenomena not explained by the exercise, one should get a hint as to what kind of abstraction process needs to be added. A conclusion of the present study is that no clear examples of such residual phenomena have been uncovered so far, even among very likely candidates, and so their number may be surprisingly small (1986, pp. 422-423).

152



Does this simulation demonstrate that Hintzman's memory model is correct? Though the results count as support for the model, which was not developed with priming or category accessibility effects in mind, that is not the real purpose of this paper. Hintzman's model has many omissions, including procedures for encoding, inference, and the like, which can also change with experience, independent of the contents of declarative memory (Smith & Branscombe, 1996). The model may well prove adequate or be falsified by experimental data gathered in different paradigms. Regardless of its eventual fate, however, the simulations reported here establish one conclusion that is *not* subject to empirical falsification: That is that an exemplar-only memory model can reproduce major features of social category-accessibility effects. The results reported here demonstrate that one member of this class of models has this ability. The point of this conclusion is that category accessibility effects can no longer be taken as evidence for the use of abstract schemas or constructs in social perception. Logically, this resembles Hintzman's demonstration that the exemplar-only memory model can reproduce effects that have been taken as evidence for prototypes in category learning. It may still be true that we have schemas or abstract constructs—in fact, I suspect that we do—but category accessibility effects do not prove it. As Hintzman notes, one benefit of this simulation will be to force more theoretical specificity concerning the role of schemas in social cognition.

153



Does this simulation demonstrate that Hintzman's memory model is correct? Though the results count as support for the model, which was not developed with priming or category accessibility effects in mind, that is not the real purpose of this paper. Hintzman's model has many omissions, including procedures for encoding, inference, and the like, which can also change with experience, independent of the contents of declarative memory (Smith & Branscombe, 1996). The model may well prove adequate or be falsified by experimental data gathered in different paradigms. Regardless of its eventual fate, however, the simulations reported here establish one conclusion that is *not* subject to empirical falsification: That is that an exemplar-only memory model can reproduce major features of social category-accessibility effects. The results reported here demonstrate that one member of this class of models has this ability. The point of this conclusion is that category accessibility effects can no longer be taken as evidence for the use of abstract schemas or constructs in social perception. Logically, this resembles Hintzman's demonstration that the exemplar-only memory model can reproduce effects that have been taken as evidence for prototypes in category learning. It may still be true that we have schemas or abstract constructs—in fact, I suspect that we do—but category accessibility effects do not prove it. As Hintzman notes, one benefit of this simulation will be to force more theoretical specificity concerning the role of schemas in social cognition.

154



Does this simulation demonstrate that Hintzman's memory model is correct? Though the results count as support for the model, which was not developed with priming or category accessibility effects in mind, that is not the real purpose of this paper. Hintzman's model has many omissions, including procedures for encoding, inference, and the like, which can also change with experience, independent of the contents of declarative memory (Smith & Branscombe, 1996). The model may well prove adequate or be falsified by experimental data gathered in different paradigms. Regardless of its eventual fate, however, the simulations reported here establish one conclusion that is *not* subject to empirical falsification: That is that an exemplar-only memory model can reproduce major features of social category-accessibility effects. The results reported here demonstrate that one member of this class of models has this ability. The point of this conclusion is that category accessibility effects can no longer be taken as evidence for the use of abstract schemas or constructs in social perception. Logically, this resembles Hintzman's demonstration that the exemplar-only memory model can reproduce effects that have been taken as evidence for prototypes in category learning. It may still be true that we have schemas or abstract constructs—in fact, I suspect that we do—but category accessibility effects do not prove it. As Hintzman notes, one benefit of this simulation will be to force more theoretical specificity concerning the role of schemas in social cognition.

155



Does this simulation demonstrate that Hintzman's memory model is correct? Though the results count as support for the model, which was not developed with priming or category accessibility effects in mind, that is not the real purpose of this paper. Hintzman's model has many omissions, including procedures for encoding, inference, and the like, which can also change with experience, independent of the contents of declarative memory (Smith & Branscombe, 1996). The model may well prove adequate or be falsified by experimental data gathered in different paradigms. Regardless of its eventual fate, however, the simulations reported here establish one conclusion that is *not* subject to empirical falsification: That is that an exemplar-only memory model can reproduce major features of social category-accessibility effects. The results reported here demonstrate that one member of this class of models has this ability. The point of this conclusion is that category accessibility effects can no longer be taken as evidence for the use of abstract schemas or constructs in social perception. Logically, this resembles Hintzman's demonstration that the exemplar-only memory model can reproduce effects that have been taken as evidence for prototypes in category learning. It may still be true that we have schemas or abstract constructs—in fact, I suspect that we do—but category accessibility effects do not prove it. As Hintzman notes, one benefit of this simulation will be to force more theoretical specificity concerning the role of schemas in social cognition.

156



Does this simulation demonstrate that Hintzman's memory model is correct? Though the results count as support for the model, which was not developed with priming or category accessibility effects in mind, that is not the real purpose of this paper. Hintzman's model has many omissions, including procedures for encoding, inference, and the like, which can also change with experience, independent of the contents of declarative memory (Smith & Branscombe, 1996). The model may well prove adequate or be falsified by experimental data gathered in different paradigms. Regardless of its eventual fate, however, the simulations reported here establish one conclusion that is *not* subject to empirical falsification: That is that an exemplar-only memory model can reproduce major features of social category-accessibility effects. The results reported here demonstrate that one member of this class of models has this ability. The point of this conclusion is that category accessibility effects can no longer be taken as evidence for the use of abstract schemas or constructs in social perception. Logically, this resembles Hintzman's demonstration that the exemplar-only memory model can reproduce effects that have been taken as evidence for prototypes in category learning. It may still be true that we have schemas or abstract constructs—in fact, I suspect that we do—but category accessibility effects do not prove it. As Hintzman notes, one benefit of this simulation will be to force more theoretical specificity concerning the role of schemas in social cognition.

157



Does this simulation demonstrate that Hintzman's memory model is correct? Though the results count as support for the model, which was not developed with priming or category accessibility effects in mind, that is not the real purpose of this paper. Hintzman's model has many omissions, including procedures for encoding, inference, and the like, which can also change with experience, independent of the contents of declarative memory (Smith & Branscombe, 1996). The model may well prove adequate or be falsified by experimental data gathered in different paradigms. Regardless of its eventual fate, however, the simulations reported here establish one conclusion that is *not* subject to empirical falsification: That is that an exemplar-only memory model can reproduce major features of social category-accessibility effects. The results reported here demonstrate that one member of this class of models has this ability. The point of this conclusion is that category accessibility effects can no longer be taken as evidence for the use of abstract schemas or constructs in social perception. Logically, this resembles Hintzman's demonstration that the exemplar-only memory model can reproduce effects that have been taken as evidence for prototypes in category learning. It may still be true that we have schemas or abstract constructs—in fact, I suspect that we do—but category accessibility effects do not prove it. As Hintzman notes, one benefit of this simulation will be to force more theoretical specificity concerning the role of schemas in social cognition.





END

11-87

DTIC