

Multiple Exemplar Training: Some Strengths and Limitations

Per Holth¹ 

Published online: 10 February 2017

© Association for Behavior Analysis International 2017

Abstract In successful multiple exemplar training, a set of exemplars that sample the range of stimulus and response topographies is trained, and the full range of performances results. Examples abound in experimental psychology and include abstraction and concept learning, responding to relations, identity matching, rule following, behavioral variability, responding to wh-questions, describing past events, learning sets, and continuous repertoires. Thus, behavior analysts often allude to a history of multiple exemplar training to account for different generalized performances. It is easy to see how a strategy of multiple exemplar training can work with many types of performances, even performances that involve relations between objects or events. However, there appear to be at least two exceptions, where direct multiple exemplar training does not work well: (1) when there are no physical dimensions at all along which generalized performances can emerge, and (2) when the relation between a stimulus and an effective response is complex. Interpretation of the latter type of cases in terms of mediated generalization is outlined and discussed. An experimental and conceptual research program should produce an account of the general limits of multiple exemplar training, and guidelines for the most effective training for generalized skills.

Keywords Multiple exemplar training · Generalized skills · Problem solving · Mediated generalization

Multiple Exemplar Training: Some Strengths and Limitations

Teaching always aims to establish more than can be directly taught. As Lovaas (1981) pointed out, generalization is a critical aspect of successful teaching. No teacher can

✉ Per Holth
per.holth@hioa.no

¹ Oslo and Akershus University College of Applied Sciences, Oslo, Norway

establish all skills in all situations. Yet, according to Baer (1981), to teach only one typical example of a skill, and then expect the students to make the necessary generalizations by themselves, may be the most common of all teacher mistakes. In their influential paper titled “An implicit technology of generalization,” Stokes and Baer (1977) described several training techniques specifically aimed at the programming of generalization, including *Train-and-Hope*. Some performances may generalize on dimensions along which no training with multiple exemplars has occurred. For example, Stokes and Baer mentioned generalization of negotiation behaviors across settings (Kifer, Lewis, Green, & Phillips, 1974). In fact, Stokes and Baer reported that many of the studies they reviewed and categorized as *Train-and-Hope* had shown successful generalization without any additional effort to that effect, such as multiple exemplar training along the relevant dimension(s).

However, when all one hoped for did not emerge, that would set the stage for other training strategies designed to assess and promote generalized performances. Lovaas (1981) reminded his readers of what he called “the basic rule about stimulus generalization: if you don’t get it, build it” (p. 110). Stokes and Baer (1977) had mentioned different techniques, such as *Sequential modification* and *Train sufficient exemplars*. The latter included sufficient exemplars of stimuli as well as responses, but Stokes and Baer did not define “exemplars” or specify what makes different instances exemplars of the same “phenomenon.” Instead, they wrote:

The optimal combination of sufficient exemplars and sufficient diversity to yield the most valuable generalization is critically in need of analysis. Is the best procedure to train many exemplars with little diversity at the outset, and then expand the diversity to include dimensions of the desired generalization? Or is it a more productive endeavor to train fewer exemplars that represent a greater diversity, and persist in the training until generalization emerges? (Stokes & Baer, 1977, p. 357).

Authors have sometimes expressed concern over these broad usages of the term “generalization,” pointing out that it discourages more detailed analyses of different basic behavioral processes involved, and may contribute to an applied field that is “more in the bag-of-tricks style than in a behavior-analytic style” (Johnston, 1979, p. 3).

My main concerns here are to (1) briefly review the empirical and theoretical basis for multiple exemplar training, (2) describe a way to deal with the concepts of classes and exemplars, (3) recount some principled limitations to what can be taught directly through multiple exemplars, and (4) suggest alternative strategies for generalized performances that lie beyond those limits.

The Empirical and Logical Basis of Multiple Exemplar Training

Hull’s (1920) hypothesis of summation provided a hypothetical-deductive background to suggest the usefulness of training multiple exemplars. According to Hull, if a particular response was reinforced in the presence of two stimuli that varied along

some dimension, both would add to the response strength in the presence of an intermediate stimulus.

Although empirical studies (e.g., Kalish & Guttman, 1959) did not support Hull's summation hypothesis, standard generalization gradients also suggested that responses could be brought effectively under control of a wider specter of a stimulus dimension by being reinforced in the presence of stimuli that varied along those dimensions (e.g., Guttman & Kalish, 1956; Hanson, 1961; Kalish & Haber, 1963). More direct evidence came from experiments in which behavior was explicitly reinforced in the presence of stimuli that varied along some specific dimension. For example, Kalish and Guttman (1959) reinforced pigeons' pecking in the presence of monochromatic light of two different wave lengths and found bimodal gradients with clear peaks of response rates on the stimulus values that had been present during reinforcement.

In sum, if a particular response is reinforced in the presence of a stimulus with particular values on some physical dimensions, and not in its absence, responses will be most likely to occur in the presence of stimuli with those same values, and gradually less likely in the presence of stimuli the more they differ from the original stimulus. Accordingly, reinforcement in the presence of stimuli that vary on the property will make responding less vulnerable to those stimulus changes.

History of Multiple Exemplar Training

A PsycINFO search for publications with “exemplar training” in either the title or in the abstract (performed on December 22nd, 2016) identified 86 publications, starting with 3 dissertation abstracts from 1985 to 1987. The first journal article to use the phrase “multiple exemplar training” (MET) appeared in 1989 and was concerned with multiple exemplars of self-instruction (Hughes & Rusch, 1989). Obviously, the use of MET was widespread long before the label was introduced. It must have played a role in instruction long before there was a system of formal education. Within experimental psychology, there is a record of it from its beginnings. For example, Thorndike (1911/2000) noted that,

Previous experience makes a difference in the quickness with which the cat forms the associations. After getting out of six or eight boxes by different sorts of acts the cat's general tendency to claw at loose objects within the box is strengthened and its tendency to squeeze through holes and bite bars is weakened; accordingly it will learn associations along the general line of the old more quickly (p. 48).

Starting with Hull (1920), training with multiple exemplars has been shown to produce the formation of what has been termed “perceptual classes” (e.g., Fields, Reeve, Matneja, Varelas, Belanich, Fitzer, & Shamoun, 2002) or “sensory-feature concepts” (e.g., Engelmann & Carnine, 1982) with many different types of stimuli in typically developing children and adults, in children with autism, and in pigeons (Herrnstein, Loveland, & Cable, 1976; Malott & Siddall, 1972; Young, Krantz, McClannahan, & Poulson, 1994). Moreover, Lovaas (1981) described how to use multiple exemplar training in order to foster more generalized skills. Such training

involves different responses, as well as using different training stimuli, and different contexts, including different places, and different people present during training.

The term multiple exemplar *instruction* (MEI) is sometimes used instead (e.g., Greer, Stolfi, Chavez-Brown, & Rivera-Valdes, 2005). Greer and Ross (2008) distinguished between two types of applications of multiple exemplar instruction: The first type is concerned with basic stimulus control through differential reinforcement in the presence of multiple stimulus exemplars with the appropriate abstract properties. The second type is concerned, instead, with bringing previously independent responses under control of a single stimulus. A more fine-grained analysis can be made by considering what constitutes an exemplar in each type of multiple exemplar training. In any case, three basic questions concern (1) what constitutes an exemplar, (2) what distinguishes different exemplars, and (3) what makes them exemplars of the same “thing” or class of phenomena. From a more practical point of view, on what basis and to what extent can we predict whether training across multiple exemplars will lead to generalized but not overgeneralized skills? It is, perhaps, comfortable to insist that these are empirical questions, but at some point, we should be able to say more than that.

Drawing the “Natural Lines of Fracture”

No functional relation between behavior and environmental events can be observed upon any single occurrence. Thus, Skinner (1935) showed that it was necessary to consider classes of stimuli and classes of responses. The class concept requires some defining property that allows for the determination of class membership. In specifying such classes, we need to take into account what Skinner called “the natural lines of fracture along which behavior and environment actually break” (Skinner, 1938, p. 33). Furthermore, as pointed out by Skinner (1969), although an operant class is primarily defined by its function, so that the topography of different instances may vary, some restriction on this variability is necessary to make possible the identification of instances. This latter point is important, yet totally disregarded when behavior analysts use such expressions as “purely functional classes,” or “reinforcement of novel behavior,” with no specifications of what counts as an instance, nor even where a response starts and where it stops.

In accord with Skinner’s previous work, Catania (1973) argued that the concept of the operant grew out of a correlation between two response classes, one descriptive and one functional. The descriptive operant is the class of responses for which consequences are arranged, and the functional operant is the class generated by that contingency. While any instance of the descriptive operant class can be identified when it occurs, we can only infer that particular instances are also members of the functional class. A functional class involves a controlling relation, and controlling relations are never directly observable (cf., Sidman, 1979). Several manipulations, and observations of instances and non-instances may be required for the identification of controlling relations, and even when a rat “lever presses for food” in an experimental chamber, any particular instance of lever pressing might occur “for other reasons.”

A descriptive operant class sometimes involves only response-descriptive features. In discriminated operants, however, the descriptive classes also involve stimuli or stimulus properties in the presence of which responses are followed by certain

consequences. By extending Catania's (1973) analysis of the operant class to the case of the discriminated operant, it is clear that the class of antecedents called discriminative stimuli (S^D) similarly requires a correlation between a descriptive and a functional class of " S^D s." Evidence of S^D control, therefore, requires a demonstration of a correlation between the stimulus class in the presence of which particular responses are reinforced and the stimulus class in the presence of which those particular responses subsequently are more likely to recur. Moreover, a specific class of S^D s extends as far, and only as far, as those instances in the presence of which those particular responses are more likely to recur.

In his effort to delineate "experimentally true" classes, Skinner (1938) mentioned the example of popular terms—which we often accept at face value without really considering whether they, in fact, refer to reliable unitary phenomena. Scientific terms must correspond to experimentally real concepts. A case in point is the generalized class called imitation. It is easy to think that a class of responses referred to as imitation almost automatically functions as a unitary phenomenon. However, Poulson and colleagues (Poulson & Kymissis, 1996; Poulson, Kyparissos, Andreatos, Kymissis, & Parnes, 2002) showed that imitation training does not automatically produce a generalized class corresponding to the generalized performances suggested by the concept of imitation. Rather, they found that when they trained normally developing infants sequentially to imitate (1) motor-with-toy movements, (2) motor-without-toy movements, and (3) vocal responses, imitation generalized within each of these classes, but not across classes. Thus, although the "popular term" imitation suggests a larger class, the actual "lines of fracture" demonstrated three different functional sub-classes. In order to answer the initial question raised by Stokes and Baer (1977) regarding what might be the most effective balance between a sufficient number and diversity of exemplars in order to obtain optimal generalization (p. 357), we need to know more about how, and how fast, functional classes come to correspond to the descriptive classes of our interventions.

Multiple Exemplar Training and Relational Frame Theory

As Skinner (1953) noted, a natural-science treatment of stimulus generalization based on a relation is unproblematic as long as the relation is identifiable in physical terms. Otherwise, alternative explanations are needed, such as an account in terms of mediating behavior. In Relational Frame Theory, however, relational responding is taken a step further, to arbitrary, not physically specifiable, relations. Thus, according to Hayes and coworkers,

... organisms could learn to respond relationally to objects where the relation is defined not by the physical properties of the object, but by some other feature of the situation.

A relational response of this kind is no longer dependent purely upon the physical properties of the relata. Rather, it is brought to bear on the stimuli encountered in the appropriate relational context: it is arbitrarily applicable. (Hayes, Fox, Gifford, Wilson, Barnes-Holmes, & Healy, 2001, p. 25).

The authors then went on to state that, although the exact answer is an empirical matter, it seems clear that relational responding involves a history of multiple exemplar training. Furthermore, the stimulus control over relational responding must have been refined by training across a variety of contexts.

Such an explanatory load on multiple exemplar training is particularly outspoken in Relational Frame Theory. For example, Hayes, Barnes-Holmes, and Roche (2001) explicitly denied that there was any need to hypothesize additional behavioral processes in addition to functional response classes established through multiple exemplars. Although the exact history involved in “transforming relational responding into an overarching arbitrarily applicable operant” may still be an empirical matter, it is *not* clear what constitutes an exemplar in the type of multiple exemplar training thought to explain how the different types of relational responding described in RFT can emerge from it.

From Multiple Exemplars to General Case

Although using multiple exemplars during teaching may often be needed to produce generalized responding, the use of multiple exemplars may not by itself suffice to produce generalized performances (e.g., Engelmann & Carnine, 1982). Sprague and Horner (1984) compared three different strategies for teaching retarded high-school students generalized skills in using vending machines. The three strategies were (1) training on only one machine, (2) training on three machines, and (3) training on three machines that sampled the whole range of variation in stimuli and responses. Whereas there was very little gain from training on three similar machines, compared with training on only one machine, the strategy of training across exemplars that actually sampled the variability in vending machines produced generalized skills in the participants of their study.

In an advanced general case model for analyzing and sequencing conditions during training (e.g., O’Neill, 1990), six steps are specified: (1) defining the Instructional Universe, i.e., the exact stimulus conditions under which specific responses should occur at the completion of training; (2) defining ranges of variability of relevant stimuli and responses; (3) selecting teaching examples that sample the full range of variations of stimulus and response properties; (4) optimizing the sequence of training examples; (5) teaching the examples, using standard state-of-the-art teaching procedures; and (6) probing with non-trained examples.

As described thus far, multiple exemplar training has been concerned with the reinforcement of different response exemplars or of reinforcement in the presence of different stimulus exemplars. Each exemplar of the second type of case consists of a class of stimuli defined by specific values on certain physical dimensions. In accord with Catania’s analysis, we may add the requirement of a correspondence between such a descriptive class and a corresponding functional class of stimuli in the presence of which responding is not significantly different. As we change the values along some physical dimension of the stimulus, and the response to it differs from the original with respect to characteristics upon which reinforcement is contingent, we have identified another exemplar in the presence of which responding may need to be directly reinforced, and so on.

Sometimes, a general case analysis (Horner, Sprague, & Wilcox, 1982) will show that there are no characteristics or dimensions that can promote a generic extension to novel cases. For example, in contrast with formal verbal operant classes (such as echoics and, to a lesser extent, textual behavior), where the establishment of minimal repertoires can spread to an extensive generalized repertoire, novel response forms in verbal operants such as tacts and intraverbals largely must be taught one by one: Having established the vocal response “cat” in the presence of a cat hardly fosters the initial emergence of “dog” in the presence of a dog or “crocodile” in the presence of a crocodile, etc. There is little or no basis for generalization across exemplars. However, even if the direct teaching of novel tacts is unlikely to gain much from multiple exemplar training, *prerequisites* for learning novel tacts from incidental teaching may very well be established that way. Each exemplar could then consist of (1) an episode where some stimulus is tacted by someone else, and (2) an occasion for the tact to be emitted and reinforced. Different sequences of combinations of (1) and (2) would then constitute different exemplars, as in what has been called “naming” (Greer et al., 2005; Horne & Lowe, 1996). Hence, there seem to be different types of multiple exemplar training, and a general case analysis of the concept of multiple exemplar training itself may be useful.

Examples of Skills that Require Multiple Exemplar Training

Generally, in multiple exemplar training following the explicit training to criterion with a first exemplar, a descriptive second exemplar is tested. To the extent that appropriate responding does not occur, its status as a second functional exemplar is confirmed. Following training with the second exemplar to criterion, a mixed training of the two exemplars to criterion is usually conducted, and followed by the testing of yet another descriptive exemplar, and so on. If at some point proposed novel exemplars evoke the appropriate response in the absence of direct training, this is evidently no longer a functionally distinct exemplar. Within a specific Instructional Universe, multiple exemplar training can be considered complete when that universe has run out of such functionally distinct exemplars. In the following, I will describe several different examples of skills that seem to require multiple exemplar training.

Abstraction and Concept Learning

Abstraction is characterized by a common response to stimulus exemplars defined by a specific property, or value on some physical dimension, but which may vary with respect to all other properties. The stimuli can be considered as functionally different exemplars to the extent that simple stimulus generalization does not occur from one to another. In accord with Catania’s (1973) analysis, we may say that an exemplar is characterized by the requirement of a correlation between such a descriptive class and a corresponding functional class of stimuli in the presence of which responding is not significantly different. As we change the values along some physical dimension of the stimulus, and the response to it differs from the original with respect to characteristics upon which reinforcement is contingent, we have identified another exemplar in the presence of which responding may need to

be directly reinforced, and so on. The point on the stimulus dimension where a particular response stops occurring marks the border of an abstract property. In a slightly more advanced set of stimulus exemplars, the exemplars are complex and vary along several dimensions (e.g., Herrnstein & Loveland, 1964). Learning based on such stimulus sets is usually characterized as concept learning. As with the different exemplars in abstraction, different exemplars of a concept may be defined by the lack of transfer from the direct reinforcement of a particular response in the presence of previous stimulus exemplars.

With appropriate multiple exemplar training with differential reinforcement, abstraction and *concept learning* can go far beyond the exemplars that are directly taught. The main limitations to abstraction and concept learning lie in the distinctions made by the environment, including the verbal community, and in the perceptual makeup of the organism. We can go beyond our perceptual limitations only through certain problem solving tactics, such as the development and use of measurement instruments, which translate non-discriminable values into discriminable ones. Microscopes, telescopes, Geiger counters, micro scales, UV meters, and so on, enable the control by abstract stimulus properties that would not control behavior effectively at all without such instruments.

Relations Between Stimuli

In another example of multiple exemplar training, each exemplar consists not of particular stimuli, but of *relations between stimuli*. Presumably, we can agree with Skinner (1953) that control by relations must be important across many environments in daily life. For example, when we move about, reinforcement is often more clearly related to relative than to absolute size. Thus, we can learn to select the biggest, the greenest, the leftmost, etc., and thus respond to a relation rather than to any absolute characteristics of stimuli. The teaching of such relational responding through multiple exemplars seems rather straightforward.

The principled limitations of responding to a relation are probably the same as those of abstraction and concept learning—as long as the relation is specifiable in physical terms. When the relation is not specifiable in physical terms, such as when the relation of concern is between different sensory fields, or when the only important relation is that two stimuli evoke the same verbal response, responding to the relation must be explained in other ways. Catania (2013) suggested that sometimes the relevant dimensions can only be specified by verbal description, but it is strange of him [Catania] to imagine that a verbal description can solve the problem, unless we can identify what controls the verbal description. The problem would still confront us.

Identity Matching

In identity matching, each exemplar consists of a sample stimulus and a positive comparison stimulus to which a specific response occurs, such as pointing to it. Non-exemplars are provided by one or more negative comparisons. Identity matching, and not just conditional discrimination, is evident when a subject responds consistently to the “matching” comparison stimulus across novel sets of stimuli. For identity

matching, the main limitations to multiple exemplar training are presumably the same as those common to abstraction and concept learning, and responding to a relation.

Rule Following

When a person is following rules, each exemplar consists of a verbal stimulus (an instruction). If a specific response is reinforced directly in the presence of such a stimulus, the resulting behavioral phenomenon would seem to be a simple example of a response controlled by a discriminative stimulus. However, after training with additional exemplars, if appropriate responses occur in the presence of novel verbal stimuli, this is more than control by directly established discriminative stimuli, and it seems appropriate to use such terms as rule following or instruction. If the rule following occurs much later, when the rule is no longer present, there appear to be a couple of possibilities. First, other variables that are present may induce the rule follower to repeat the rule so that the rule is, in fact, present again at this time. Second, the original verbal stimulus may have had functioning-altering effects by establishing stimuli that are now present as positive or negative discriminative stimuli or reinforcers (Schlinger & Blakely, 1987).

How does rule following spread across different rules (different responses, different S^D s, and different reinforcers and motivational operations)? It is difficult to trace the exact “development” of generalized instruction following in the world at large, but curricula for children with learning deficits typically describe multiple exemplar training in which several things are more or less systematically varied, including different verbs, prepositions, objects to respond to, occasions upon which to respond, and the context in which the instruction is presented (e.g., Lovaas, 1981; Luiselli, Russo, Christian, & Wilczynski, 2008; Maurice, Green, & Luce, 1996). Still, rule following can hardly generalize much beyond the description of responses, objects, prepositions, and motivational operations to which the individual has been taught directly to respond. One exception to this limitation may be those cases in which learning by exclusion is made possible. I will return to that in the context of “learning set” below. Another exception, to overcome the limitation, involves problem solving skills in the form of asking about the “meaning” of, or synonyms for unfamiliar words contained in the instruction. By including multiple exemplar training of such problem solving, rule following skills can extend far beyond what can otherwise be obtained through direct multiple exemplars of rules exclusively.

Lag N Reinforcement Schedule

A *Lag N schedule* of reinforcement involves the criterion that the response differs from N previous responses. Each exemplar is then characterized by the occurrence of the specific responses that constitute the last N responses. When the responses are relatively simple and the N is low, the so-called memory-based responding can develop. For example, if there are three response options and a Lag 2 schedule, and Response 1 and Response 3 have occurred most recently, then only Response 2 is eligible for reinforcement. Thus, following multiple exemplars, a generalized class of responses, consisting of responding differently from the last two responses, may develop. The different exemplars here consist of different combinations of “last two responses” from which a current response must differ in order to produce the reinforcer. A multiple

exemplar training of this sort can only work up to a certain level of complexity of the response options and up to a certain N .

When the responses are more complex or the N is much higher, “memory-based” responding is less likely, and the resulting behavior looks random or stochastic. Neuringer and colleagues (e.g., Neuringer, 2009; Page & Neuringer, 1985) have argued that the resulting behavioral variability itself should be considered as an operant. However, the resulting variability is typically restricted to the specific responses that have been followed by reinforcement, and that variability can be accounted for in terms of the cyclic reinforcement and extinction of these responses (see Holth, 2012). Thus, lag schedules simply have an inherent characteristic which ensures that no specific responses or response sequences are differentially reinforced at the cost of others in the long run. Only when the lag is very low and the response or response sequence is very simple do multiple exemplars of lag training lead to higher-order “stereotypy” in which previous responses enter into the contingency as S^A s for repeating those responses.

Responding to Wh-Questions

To a large extent, teaching responding to wh-questions (who, where, what, which, and when) seems like a very straightforward task (e.g., Jahr, 2001). At least following the reading or telling of a story, or following the direct observation of some event, responding to questions about who, where, what, which, and when, should be a matter of direct multiple exemplar training. However, responding to why questions may be somewhat more complicated, because why questions seem to require several different types of answers. If asked why someone is eating, we have learned to report some antecedent event or, rather, some feeling associated with a prior event, such as hunger associated with food deprivation, thirst associated with liquid deprivation, and so on. It is difficult to see how any number of such exemplars could generalize to answers to other categories of why questions, such as *Why do you go to school?* *Why does the wolf have such big ears?* *Why does the giraffe have such a long neck?* and *Why did the glass break?* Indeed, some why questions are answered differently in different verbal communities. For example, why questions regarding behavior require different answers in a traditional cognitive psychology community than in a behavior-analytic community, and yet other answers in a neurology context (cf., Holth, 2013).

Describing Past Events

Presumably, an exemplar of describing past events consists of (a) an event, (b) the passage of time, and (c) an occasion on which the reporting of past events is appropriate, such as someone asking “What did you do last night?” or “What did you have for breakfast?” which then controls a response that would have been characterized as a standard tact had it occurred in the presence of the past event. Skinner discussed such examples:

... when the child says *There was an elephant at the zoo*, he appears to be reacting to his past history rather than merely profiting from it. This is a verbal achievement brought about by a community which continually asks the child such questions as *Was there an elephant at the zoo?* The answer must be

understood as a response to current stimuli, including events within the speaker himself generated by the question, in combination with a history of earlier conditioning. (Skinner, 1953, p. 178).

The question, then, is what, specifically, constitutes a relevant history of early conditioning. Skinner mentioned that the verbal community continually asks questions about what has happened. Presumably, this adds up to a kind of multiple exemplar training, but how does “responding to past events” generalize across questions concerning different persons, different sense modalities, and different times, such as *What did you do? What did Anna do? What did you see? What did you hear? What did you smell? And What did you taste? What did you have for breakfast today? What did you have yesterday? What did you do Monday evening?* If starting from scratch with a child who completely lacked the relevant repertoires, how would a good behavior analyst go about establishing these skills most effectively? Presumably, in short, the strategy would consist of an initial multiple exemplar training similar to the one prescribed for concept learning and for answering wh-questions here and now. Next, one might gradually extend the time from the actual event till the questions about the past event are asked. Under such contingencies, “memory skills” may “develop” even if none are explicitly taught. Examples of relevant skills have been described by Palmer (1991) under the heading, “Memory as Problem Solving.”

Learning Sets

Yet another type of exemplars seems to be involved in phenomena referred to as *learning set*. As originally described by Harlow (e.g., 1949), on each trial, two different objects are presented to a monkey. The monkey is food deprived, and food is always accessible only by picking up one of the objects. Over successive trials with the same two objects, the monkey ends up consistently picking up the one under which food is hidden. After similar exposure to hundreds of object pairs, the monkeys eventually learn from the single first trial so that from the second trial on, they consistently choose the object that covers the food. In this case, an exemplar consists of one stimulus that is systematically correlated with reinforcement contingent on a specific type of response and one that is not, until these stimuli are established as an S^D and an S^Δ , respectively, for such responses. Over successive repetitions with the same exemplar, the functions of stimuli are gradually established. Across successive presentations of multiple exemplars, this function-altering effect speeds up, and eventually, a single trial with a novel exemplar is sufficient to reliably alter the function of the stimuli involved. In Harlow’s experiment, the function-establishing stimulus was the consequence of the first response but, presumably, it does not have to be response-contingent. The intermixing of different stimulus exemplars is largely irrelevant in this type of multiple exemplar training, because the success criterion lies not in the number of discriminations mastered, but in the speed with which new ones are formed.

The *mutual entailment* (e.g., if $A > B$, then it follows that $B < A$) described in relational frame theory (e.g., Hayes et al., 2001) may typically result from multiple exemplar training of this type. Each exemplar then consists of a pair of objects, as in the relational case (2) above, but each exemplar also includes a second trial, in

which the relation specified in the first trial (e.g., $A > B$) and some conditional stimulus (e.g., “How is B related to A?”) jointly controls a different response (e.g., “It’s smaller”). A similar “opposite relation” is implicit in Harlow’s learning set: If the object picked up on the first trial was correct (i.e., picking it up produced the reinforcer), then the other object is incorrect and will be consistently avoided on a second trial. If the object picked up on the first trial was incorrect (i.e., picking it up was not reinforced), then the other object is correct and will be consistently chosen on a second trial. Hence, the mutual entailment seems explainable by what is appropriately called multiple response-exemplar training, in which two or more responses are established under the control of the same stimulus material. For example, in the presence of a big A and a smaller B, both “A is bigger than B” and “B is smaller than A” are directly taught as different response exemplars under control of the same stimulus material.

Yet another learning-set type of multiple exemplar training brings responses under joint stimulus control and, thus, produce what is described as “generative” phenomena, such as naming (e.g., Greer et al., 2005; Greer, Stolfi, & Pistoljevic, 2007; Horne & Lowe, 1996) and the emergence of untaught responses across different verbal functions, such as from mands to tacts and vice versa (e.g., Greer & Ross, 2008). In these cases, each exemplar consists of a sequence of trials during which additional stimuli, such as “What is that?” and “Point to X,” correlate with the reinforcement of different responses controlled by the same stimulus material. Eventually, following training with multiple exemplars, a single exposure to a novel vocal stimulus in the presence of a novel stimulus material is sufficient for the additional stimuli to alter the S^D function of the same stimulus material differently. The additional stimuli may be designated conditional stimuli (S^C s). Depending on which S^C is present on a given occasion, the same stimulus material may serve as an S^D for a tact, as an S^D for pointing to it, or otherwise selecting it.

Continuous Repertoires

In some relational cases, neither the stimuli that are present nor the responses remain the same across exemplars, but effective responses vary with variations in some property of a stimulus. This is different from just stimulus or response generalization. Such cases have sometimes been treated under the heading *continuous repertoires* (e.g., Wildemann & Holland, 1972). In a simple example, let us say we have an array of seven light sources arranged on a horizontal line from left to right. Below that line, we have a corresponding array of seven levers from left to right. Then, in the presence of the leftmost light (Light 1), we reinforce a rat’s responses on the leftmost lever (Lever 1). Then we may ask what happens if we, instead, turn on Light 5. Surely, at this point, our rat is likely to continue pressing Lever 1. A second exemplar could then consist of the presentation of Light 5 in the presence of which presses on Lever 5 will be reinforced. Following the mixed training of these, what will our rat do if we turn on Light 3, or Light 7? Each light may be characterized as a separate descriptive exemplar, but following additional training, at some point, each novel light is no longer accompanied functionally distinct classes. Instead, one higher-order class may be described as lever-press

positions controlled by light positions. This result would be a simple example of a continuous repertoire.

As in the first type of multiple exemplar training mentioned above (i.e., abstraction and concept learning), the relevant properties of stimuli in continuous repertoires can be multidimensional, as for example in imitation: All sorts of dimensions in the behavior of a model can be reflected in the behavior of the imitator.

Clearly, there must also be limits to what kinds of continuous repertoires can be established through direct training with multiple exemplars. At some level of complexity, the S^D ceases to function as such with respect to the effective response. Hence, for example, our imitation of a long sentence in an unfamiliar language is likely to be quite imperfect. When, in spite of this, some apparent functional relation between an antecedent S and the R still remains, the relation must be bridged or mediated by additional events.

Mediated Generalization

In the example above, with response positions controlled by light positions, a child, or even a rat, could achieve this by observing the light and then gaze vertically down and then stop when arriving at the corresponding response position. Such a simple strategy would not be very useful across different, more complex tasks (for example, one in which the stimulus and response dimensions rotated relative to each other or one that required the solving of mathematical problems). If, however, instead the participant emits a differential response, such as uttering the position number on a stimulus array (possibly as an end point of counting positions from left to right) and repeats that number until the same response topography is jointly controlled by the corresponding number on the response array, such a strategy would be very useful across a range of different kinds of problems (e.g., Lowenkron, 1991, 1996, 1998).

When the complexity increases so that the relation between a preceding stimulus and the appropriate response becomes difficult to describe at all in physical terms, the less likely it seems that the appropriate novel responding results automatically from a particular identifiable contingency. For instance, no number of exemplars with the multiplication of three- and four-digit numbers would suffice to establish a repertoire of generalized correct responses to novel exemplars in the absence of some kind of precurent problem solving skills that allow for what Stokes and Baer (1977) called “mediated generalization” (p. 131).

Stokes and Baer (1977) described mediated generalization as a technique in which responses that are taught as part of a particular training program become functional under additional circumstances:

More precisely, mediated generalization involves the strategic use of person-transported antecedent stimuli and responses as controlling variables to enhance performance across circumstances and across time. That is, a relevant stimulus or behavior is incorporated as part of initial learning, which, when produced across diverse (generalized) contexts by the person, occasions the

performance of the relevant behaviors. It is the individual's liaison or transfer of the controlling stimuli across situations that provides the mediating conditions to enhance performance across the generalization dimensions. (Stokes & Baer, 2003, pp. 131–132).

Much of what is explicitly taught in schools constitutes techniques for mediated generalization. For example, children learn specific strategies to add numbers, to subtract, to multiply, and to divide. Learning to do these things on the calculator is just a special case. Schools do not waste time on direct multiple exemplar training where the student is expected to solve new problems correctly simply as a function of having learned to respond correctly to a number of different math assignments. As a function of being taught to respond in accordance with a specific *strategy* on multiple problems, the student is expected to respond in accordance with the same strategy to solve novel problem exemplars. Common examples of mediators are rules, or self-instructions, but less explicit, and perhaps nonverbal, strategies can probably also function similarly.

When students learn to solve a certain type of equation problems, it is typically through multiple exemplars relying on the same general pattern. Thus, it is possible to learn certain strategies that speed up the solving of specific types of problems: One procedure for multiplying, another for dividing, yet other procedures for solving equations, and so on. A specific problem-solving skill is established when a specific strategy is successfully applied to novel exemplars. To what extent is it also possible, through multiple exemplars of types of problems, to learn more general problem solving skills, so that problem solving in general can become a skill? A kind of second-order problem solving may be needed when the problem confronted does not function as an S^D for engaging in a problem-solving strategy (precurrent behavior) that has been directly trained.

Precurrent Behavior Explicitly or Implicitly Taught

As Skinner (1968) pointed out, precurrent responses need not be explicitly reinforced to be maintained. It is probably sufficient that they increase the probability of reinforcement of the terminal response. However, the terminal response or “solution” may not necessarily be explicitly reinforced either. Sometimes, it seems sufficient to “know that one is right.” When precurrent behavior is not terminated by some explicit reinforcement of a successful response, we must point to other variables to account not only for why the precurrent behavior starts but also for why it stops, when we “know we are right.” To say that we “know something” is, of course, verbal behavior, and like other verbal behavior, it may be controlled by a plethora of different independent variables. However, one type of control that may be characteristic of the kinds of situations in which we claim that we “know we’re right” is when our solution to a problem is jointly controlled by different variables: Our numbers in a Sudoku fit vertically as well as horizontally, our numbers in a multiplication problem can be reversed through division, and our puzzle-board pieces fit with adjoining pieces as well as with the larger picture of the board, and so on.

Conclusion

Although multiple exemplar training may effectively establish a number of different generalized skills, for some performances, it seems clear that no number of multiple exemplars of direct training can suffice to establish the general skill. In a behavior analysis that aims to account for moment-to-moment changes in behavior based on a minimal number of basic, or general, principles, it seems that mediated generalization must play an important role in an account of many complex cases of problem-solving behavior. Several basic facts about such mediators are still to be investigated, such as (1) what characterizes precurent behavior that most effectively produces generalized correct responses to an array of problems; (2) what are the most effective strategies, and sequencing, for teaching such mediators; and (3) when should such mediating behavior be taught explicitly, and when should it be left to natural contingencies without explicit training? These interesting themes should definitely not be left to cognitive psychology, which will not reveal the kinds of independent variables that practitioners will find most useful.

As Stokes and Baer (2003) suggested, even if mediated generalization is still underdeveloped at present, it may have a great potential as a spearhead for producing behavioral change. Stokes and Baer characterized it as an “unfinished portrait” in the analysis of strategies aiming to promote generalization. The inclusion of mediated generalization in a comprehensive account of multiple exemplar training will be necessary in order to develop our guidelines for the most effective training for generalized skills.

Compliance with Ethical Standards

Conflict of Interest The author declares that he has no conflict of interest.

References

- Baer, D. M. (1981). *How to plan for generalization*. Austin: Pro-Ed.
- Catania, A. C. (1973). The concept of the operant in the analysis of behavior. *Behaviorism*, 1, 103–116.
- Catania, A. C. (2013). *Learning* (5th ed.). Cornwall on Hudson: Sloan Publishing.
- Engelmann, S., & Carnine, C. (1982). *Theory of instruction: principles and applications*. New York: Irvington Publishers, Inc.
- Fields, L., Reeve, K. F., Matreja, P., Varelas, A., Belanich, J., Fitzer, A., et al. (2002). The formation of a generalized categorization repertoire: effects of training with multiple domains, samples, and comparisons. *Journal of the Experimental Analysis of Behavior*, 78, 291–313.
- Greer, R. D., & Ross, D. E. (2008). *Verbal behavior analysis: inducing and expanding new verbal capabilities in children with language delays*. Boston: Allyn & Bacon.
- Greer, R. D., Stolfi, L., Chavez-Brown, M., & Rivera-Valdes, C. (2005). The emergence of the listener to speaker component of naming in children as a function of multiple exemplar instruction. *The Analysis of Verbal Behavior*, 21, 123–134.
- Greer, R. D., Stolfi, L., & Pistoljevic, N. (2007). Emergence of naming in preschoolers: a comparison of multiple and single exemplar instruction. *European Journal of Behavior Analysis*, 8, 109–131.
- Guttman, N., & Kalish, H. I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology*, 51, 79–88.

- Hanson, H. M. (1961). Stimulus generalization following three-stimulus discrimination training. *Journal of Comparative and Physiological Psychology*, 54(2), 181–185.
- Harlow, H. F. (1949). The formation of learning sets. *Psychological Review*, 56(1), 51–65. doi:10.1037/h0062474.
- Hayes, S. C., Barnes-Holmes, D., & Roche, B. (2001). Relational frame theory: a précis. In S. C. Hayes, D. Barnes-Holmes, & B. Roche (Eds.), *Relational frame theory: a post-Skinnerian account of human language and cognition* (pp. 141–154). New York: Kluwer.
- Hayes, S. C., Fox, E., Gifford, E. V., Wilson, K. G., Barnes-Holmes, D., & Healy, O. (2001). Derived relational responding as learned behavior. In S. C. Hayes, D. Barnes-Holmes, & B. Roche (Eds.), *Relational frame theory: a post-Skinnerian account of human language and cognition* (pp. 21–49). New York: Kluwer.
- Herrnstein, R. J., & Loveland, D. H. (1964). Complex visual concept in the pigeon. *Science*, 146, 549–551.
- Herrnstein, R. J., Loveland, D. H., & Cable, C. (1976). Natural concepts in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 4, 285–301.
- Holth, P. (2012). Variability as an operant? *Behavior Analyst*, 35(2), 243–248.
- Holth, P. (2013). Different sciences as answers to different why questions. *European Journal of Behavior Analysis*, 14, 165–170.
- Horne, P. J., & Lowe, C. F. (1996). On the origins of naming and other symbolic behavior. *Journal of the Experimental Analysis of Behavior*, 65, 185–241.
- Horner, R. H., Sprague, J., & Wilcox, B. (1982). General case programming for community activities. In B. Wilcox & G. T. Bellamy (Eds.), *Design of high school programs for severely handicapped students* (pp. 61–98). Baltimore: Brookes.
- Hughes, C., & Rusch, F. R. (1989). Teaching supported employees with severe mental retardation to solve problems. *Journal of Applied Behavior Analysis*, 22, 365–372. doi:10.1901/jaba.1989.22-365.
- Hull, C. L. (1920). Quantitative aspects of evolution of concepts: an experimental study. *Psychological Monographs*, 28, i–86. doi:10.1037/h0093130.
- Jahr, E. (2001). Teaching children with autism to answer novel wh-questions by utilizing a multiple exemplar strategy. *Research in Developmental Disabilities*, 22, 407–423.
- Johnston, J. M. (1979). On the relation between generalization and generality. *Behavior Analyst*, 2, 1–6.
- Kalish, H. I., & Guttman, N. (1959). Stimulus generalization after training on three stimuli: a test of the summation hypothesis. *Journal of Experimental Psychology*, 57, 268–272.
- Kalish, H. I., & Haber, A. (1963). Generalization: I. Generalization gradients from single and multiple stimulus points. II. Generalization of inhibition. *Journal of Experimental Psychology*, 65, 176–181.
- Kifer, R. E., Lewis, M. A., Green, D. R., & Phillips, E. L. (1974). Training predelinquent youths and their parents to negotiate conflict situations. *Journal of Applied Behavior Analysis*, 7, 357–364.
- Lovaas, O. I. (1981). *Teaching developmentally disabled children: the me book*. Baltimore: University Park Press.
- Lowenkron, B. (1991). Joint control and the generalization of selection-based verbal behavior. *The Analysis of Verbal Behavior*, 9, 121–126.
- Lowenkron, B. (1996). Joint control and word-object bidirectionality. *Journal of the Experimental Analysis of Behavior*, 65, 252–255.
- Lowenkron, B. (1998). Some logical functions of joint control. *Journal of the Experimental Analysis of Behavior*, 69, 327–354.
- Luiselli, J. K., Russo, D. C., Christian, W. P., & Wilczynski, S. M. (Eds.). (2008). *Effective practices for children with autism*. New York: Oxford University Press.
- Malott, R., & Siddall, J. W. (1972). Acquisition of the people concept in the pigeon. *Psychological Reports*, 31, 3–13.
- Maurice, C., Green, G., & Luce, S. (Eds.). (1996). *Behavioral intervention for young children with autism: a manual for parents and professionals*. Austin: Pro-Ed.
- Neuringer, A. (2009). Operant variability and the power of reinforcement. *The Behavior Analyst Today*, 10(2), 319–343.
- O'Neill, R. E. (1990). Establishing verbal repertoires: toward the application of general case analysis and programming. *The Analysis of Verbal Behavior*, 8, 113–126.
- Page, S., & Neuringer, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Processes*, 11, 429–452.
- Palmer, D. C. (1991). A behavioral interpretation of memory. In L. J. Hayes & P. N. Chase (Eds.), *Dialogues on verbal behavior* (pp. 261–279). Reno: Context Press.

- Poulson, C. L., & Kymissis, E. (1996). The operant language-acquisition paradigm and its empirical support. In S. W. Bijou & E. Ribes (Eds.), *New directions in behavior development* (pp. 73–89). Reno: Context Press.
- Poulson, C. L., Kyparissos, N., Andreatos, M., Kymissis, E., & Parnes, M. (2002). Generalized imitation within three response classes in typically developing infants. *Journal of Experimental Child Psychology*, *81*, 341–357. doi:10.1006/jecp.2002.2661.
- Schlinger, H., & Blakely, E. (1987). Function-altering effects of contingency-specifying stimuli. *The Behavior Analyst*, *10*(1), 41–45.
- Sidman, M. (1979). Remarks. *Behaviorism*, *7*(2), 123–126.
- Skinner, B. F. (1935). The generic nature of the concepts of stimulus and response. *The Journal of General Psychology*, *12*(1), 40–65.
- Skinner, B. F. (1938). *The behavior of organisms: an experimental analysis*. New York: Appleton-Century.
- Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.
- Skinner, B. F. (1968). *The technology of teaching*. New York: Appleton.
- Skinner, B. F. (1969). *Contingencies of reinforcement: a theoretical analysis*. Englewood Cliffs: Prentice Hall.
- Sprague, J. R., & Horner, R. H. (1984). The effects of single instance, multiple instance, and general case training on generalized vending machine use by moderately and severely handicapped students. *Journal of Applied Behavior Analysis*, *17*, 273–278.
- Stokes, T. F., & Baer, D. M. (1977). An implicit technology of generalization. *Journal of Applied Behavior Analysis*, *10*, 349–367.
- Stokes, T. F., & Baer, D. M. (2003). Mediated generalization: an unfinished portrait. In K. S. Budd & T. Stokes (Eds.), *A small matter of proof: the legacy of Donald M. Baer* (pp. 125–138). Reno: Context Press.
- Thorndike, E. L. (1911/2000). *Animal intelligence: experimental studies*. New Brunswick: Transaction Publishers.
- Wildemann, D. G., & Holland, J. G. (1972). Control of a continuous response dimension by a continuous stimulus dimension. *Journal of the Experimental Analysis of Behavior*, *18*, 419–434.
- Young, J. M., Krantz, P. J., McClannahan, L. E., & Poulson, C. L. (1994). Generalized imitation and response-class formation in children with autism. *Journal of Applied Behavior Analysis*, *27*, 685–697.