

THE FUNCTION ACQUISITION SPEED TEST (FAST): A BEHAVIOR ANALYTIC IMPLICIT TEST FOR ASSESSING STIMULUS RELATIONS

Anthony O'Reilly and Bryan Roche
National University of Ireland, Maynooth

Maria Ruiz
Rollins College, Florida

Ian Tyndall
Chichester University

Amanda Gavin
Teesside University

Subjects completed a baseline stimulus matching procedure designed to produce two symmetrical stimulus relations; A1–B1 and A2–B2. Using A1, B1, and two novel stimuli, subjects were then trained to produce a common key-press response for two stimuli and a second key-press response for two further stimuli across two blocks of response training. During one block, the reinforcement contingencies were consistent with baseline relations (i.e., A1 and B1 shared a response function), whereas during the other block they were not. Thirteen of 18 subjects who completed the procedure showed a response class acquisition rate differential across the two test blocks in the predicted direction. It is suggested that this procedure may serve as a behavior analytic alternative to popular implicit tests. It provides a nonrelative measure of stimulus association strength and may display superior procedural implicitness over other tests.

Key words: FAST, Function Acquisition Speed Test, Implicit Association Test, Implicit Relational Assessment Procedure, implicit cognition

There has been considerable recent interest in developing behavior analytic “implicit” tests for assessing histories of relational responding and stimulus relations generally. This interest can be traced to the finding that the stimulus relations formed during a subject’s social history may interfere with the formation of novel stimulus relations, such as equivalence classes. Specifically, in what can now be surely described as a seminal study, Watt, Keenan, Barnes, and Cairns (1991) used a simple stimulus

The authors would like to thank Lanny Fields and Dermot Barnes-Holmes for comments and suggestions during the preparation of this manuscript.

Correspondence concerning this article should be addressed to Dr. Bryan Roche, Department of Psychology, National University of Ireland, Maynooth, Co. Kildare, Ireland. E-mail: Bryan.T.Roche@nuim.ie

equivalence paradigm to identify a history of religious sectarianism among Northern Irish and English subjects. These researchers took advantage of the fact that people in Northern Ireland often respond to names as indicative of social and religious background. They suspected that well-established stimulus relations between real-word stimuli might interfere with the acquisition of relationally incongruous stimulus equivalence classes.

In the Watt et al. (1991) study, subjects were first exposed to a matching-to-sample (MTS) procedure, in which A–B relations were trained between Catholic surnames and nonsense syllables, followed by B–C training between the nonsense syllables and Protestant symbols, leading to derived A–B–C equivalence classes. During equivalence testing, English subjects correctly matched the Catholic names with the Protestant symbols (i.e., derived A–C relations), but 12 of 19 Northern Irish subjects chose a novel Protestant name in the presence of the Protestant symbols, thereby failing to demonstrate the normally predictable stimulus equivalence classes. These findings suggested that the social contingencies operating in Northern Ireland interfered with the establishment of equivalence relations in the laboratory. In simple terms, the derived equivalence testing procedure required Northern Irish subjects to juxtapose names and symbols in a manner that was counter cultural. The outcome of this study laid the groundwork for a behavior analytic test that might allow researchers to tap into stimulus relations, the derivation or discrimination of which has been supported by the appropriate contingencies but which has not in fact yet been derived or discriminated. It might also allow researchers to assess specific relations between stimuli, without alerting subjects to the point of the test. In other words, the test was shown to be capable of identifying stimulus relations implicit in the social history of a subject while at the same time circumventing the need for direct questioning and the associated demand characteristics (i.e., social desirability; see Paulhus, 2002).

The Watt et al. (1991) procedure certainly captured the attention of behavior analysts, who have used the methodology extensively. Dixon, Rehfeldt, Zlomke, and Robinson (2006) found that the Watt et al. (1991) method could be used to measure the social discrimination of Middle-Eastern people. In a study on gender identity, Moxon, Keenan, and Hine (1993) found that subjects had more difficulty forming equivalence classes when the classes included female names and stereotypic male occupations (see also Kohlenberg, Hayes, & Hayes, 1991; Roche & Barnes, 1996). In another study, Leslie et al. (1993) demonstrated that anxious subjects had difficulty in matching pleasant-state adjectives to stimuli representing threatening situations in a stimulus equivalence test. Barnes, Lawlor, Smeets, and Roche (1995) employed the Watt et al. (1991) paradigm to examine the formation of stimulus equivalence relations between positive and negative words and subjects' own names (see also Merwin & Wilson, 2005). McGlinchey, Keenan, and Dillenburger (2000) harnessed the stimulus equivalence test paradigm in developing a method to identify the presence of child sexual abuse among children. Finally, Roche, Ruiz, O'Riordan, and Hand (2005) also reported on the use of the Watt et al. (1991) paradigm in a pilot study that attempted to distinguish child sex offenders from a sample of control subjects.

The foregoing studies all provided promise that the Watt et al. (1991) method could be harnessed to produce subtle tests for social and other sensitive stimulus relations that may not be easily measurable using explicit self-report methodologies, such as paper-and-pencil tests. However, these studies all involved the use of real words as stimuli and relied upon speculation about the types of stimulus relations that already existed in the social histories of the subjects. What was needed was a dedicated laboratory controlled analysis of the Watt et al. *process*, using arbitrary stimuli whose functions were under experimental control.

The first study to attempt such an analysis was conducted by Roche, Barnes, and Smeets (1997; Experiment 2). The authors paired two nonsense syllable stimuli (A1 and C2) with sexually arousing film clips and two other stimuli (A2 and C1) with nonsexual film clips. The A1–C2 and A2–C1 relations were then tested using a matching-to-sample test.

The experimenters then trained subjects on a matching-to-sample procedure that was designed to lead to the emergence of two 3-member equivalence classes among nonsense syllable stimuli, including the conditioned sexual and nonsexual stimuli. The emergent relations were each intended to contain one conditioned sexual and one conditioned nonsexual stimulus. In other words, the emergent equivalence relations (e.g., A1–B1–C1) would be incongruous with the previously established functional classes (e.g., A1–C2). When reexposed to the matching-to-sample testing, subjects matched stimuli based on their conditioned sexual/nonsexual functions rather than forming the equivalence classes predicted by the MTS procedure. This demonstrated that the emergence of stimulus equivalence can be impeded by the existence of incongruous stimulus relations. This research, therefore, corroborated the Watt et al. (1991) suggestion that stimulus equivalence tests can be used as tests for the existence of previously established stimulus relations in socially sensitive contexts (i.e., in which direct questioning of a subject is inappropriate or ineffective; see also Hall, Mitchell, Graham, & Lavis, 2003; Tyndall, Roche, & James, 2004, 2009, for related studies).

While the Watt et al. (1991) procedure has proven to be of considerable interest to behavior analysts, it has not been adopted widely as a testing methodology in the world outside the laboratory or by psychologists in mainstream psychology, as have several popular “implicit” tests, such as the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998) and the emotional Stroop (see Williams, Mathews, & Macleod, 1996, for a review). We offer three main reasons why this is the case. First, the Watt et al. stimulus equivalence paradigm requires an understanding of the stimulus equivalence phenomenon and the conditional discrimination training methodology with which it is usually associated. Thus, it is not easily accessible to non-behavior analysts. Second, the test format is very demanding of subjects insofar as that it requires considerable attention and motivation in order to complete training and testing phases of the procedure. Finally, the test format requires considerable time compared to other popular implicit tests, such as the IAT or emotional Stroop. While the latter typically requires only a few minutes, the stimulus equivalence methodology requires approximately 20 minutes to complete.

Another potential shortcoming of the Watt et al. (1991) procedure is one that may compromise its subtlety. Specifically, in this procedure, two incompatible stimuli are presented together as sample and comparison, respectively. While a subject has no way of being certain what the purpose of the measure is, the fact remains that the stimuli whose relation is under analysis are presented simultaneously during the testing phase. Thus, procedural implicitness is compromised by the Watt et al. technique. *Procedural implicitness* refers to the degree to which the relations under analysis and the purpose of the test are discriminable by a subject. This is to be distinguished from *outcome implicitness*, which refers to the implicitness of the stimulus associations being measured (i.e., whether the subject has ever discriminated the associations in the past; see De Houwer, 2006).

An important but simple modification can retain the basic core process of the Watt et al. methodology while at the same time disguising the purpose of the test more fully. More specifically, if we attempt to establish a functional response class instead of a derived equivalence relation, an entirely novel test format presents itself. That is, instead of presenting the two stimuli of interest simultaneously to assess the probability of a matching response during an equivalence test, it is possible to present the stimuli individually on separate trials and attempt to establish a functional stimulus class containing them. This can be done by establishing distinct response functions for stimuli that are suspected of being related as a result of prior social interaction. The rate at which a subject learns to produce the common response for both of these stimuli, presented separately, can be compared to the learning rate for producing distinct responses for these two stimuli, again across separate trials. Such a procedure requires no conditional discrimination training, is not demanding on the subject (i.e., there are no relations to derive), and is fast to administer.

As an example of the foregoing, we might assess whether or not the words *mathematics* and *fun* are related for a sample of school children with difficulty learning math. We could present them with each of four stimuli, individually and in a random order, and across several re-cycled blocks of trials. The stimuli should consist of the words of interest (*mathematics* and *fun*) and also two novel and entirely unrelated words; nonsense syllables may work best. Upon the presentation of each stimulus, the subject may be required to produce one of two responses. These responses may be positional on a computer keyboard or may involve clicking on a specific discriminated response key on a computer screen. Correct responses can be easily reinforced verbally or with tokens. It is important that the words of interest share a common response, while the novel control stimuli (e.g., nonsense syllables) also share the remaining response function. The number of trials required for the subject to reach a preset fluency criterion in this block of trials represents an index of the preexisting strength of the relation between the words *mathematics* and *fun*.

We may also rerun the test block, but with the important difference that now we establish a response class for the word *mathematics* and one of the control words. Similarly, the word *fun* and the remaining control word now share a response. If the words *mathematics* and *fun* have been related in the history of the subject, then this latter task block will require a larger number of training trials than the former block described above, simply because the latter task is inconsistent with the social history of the subject (i.e., the functional class is inconsistent with the socially established functional or equivalence class containing the words *mathematics* and *fun*). A baseline rate of acquisition of response classes using novel and arbitrarily chosen stimuli may also be recorded. We may then assess the extent of the facilitating or retarding effect of preexisting stimulus relations between *mathematics* and *fun* on learning during the two test blocks. Taken together, this approach to assessing the strength of a stimulus–stimulus relation can be referred to as a Function Acquisition Speed Test (FAST).

The rationale underlying the FAST format is not unlike that used by researchers studying resistance to change effects. Nevin and Grace (2000) described this effect using the concept of behavioral momentum. They suggest that momentum describes rate of responding, which is analogous to the velocity of a physical body in motion under Newtonian physics. When high-rate stable behavior is disrupted in some way by a change in contingency, we see effects that are analogous to the inertial mass of a physical body in motion. That is, behavior (or movement) continues on the original trajectory, and current contingencies are less effective than they would be if current behavior was low in momentum. In a representative experiment, Nevin (1974, Experiment 1) trained pigeons on multiple VI 1-minute (Component 1) and VI 3-minute schedules (Component 2). Once behavior was stable, a disruptor in the form of “time out” food between components was introduced. Rates of responding were higher preceding the presentation of the disruptor, and the decrease in response rate was smaller in Component 1 compared to Component 2. In such studies, the disrupting effect is typically quantified by calculating the proportion of responding under the novel contingency compared to baseline rates. Logarithmic transformations of this proportion are also usually applied. This index is then used as an index of strength of responding given a particular stimulus, a technique which can also be usefully applied to the FAST procedure.

The FAST methodology also has much in common with the widely used IAT (see Gavin, Roche, & Ruiz, 2008; Gavin, Roche, Ruiz, Hogan, & O'Reilly, 2012; Ridgeway, Roche, Gavin, & Ruiz, 2010; Roche et al., 2005). The IAT was originally said to measure “implicit” attitudes or beliefs, while preventing subjects from exerting conscious control over the test outcomes (Banse, Seise, & Zerbes, 2001; Kim, 2003). A guiding but theory-uncommitted assumption of the IAT (see Greenwald, Nosek, Banaji, & Klauer, 2005) is that it is easier to assign a single response to two concepts if they are associated in memory than if they are unrelated.

One early IAT study (Greenwald et al., 1998) involved presenting subjects with flower names (e.g., TULIP) insect names (e.g., SPIDER), pleasant words (e.g., LOVE) or unpleasant words (e.g., UGLY) individually, on separate trials. Subjects were asked to categorize the names by means of one of two positional key presses. In the first (consistent) condition, the same response key was assigned to both flower and pleasant stimulus words, while insects and unpleasant word stimuli shared a different but common positional response. In the second (inconsistent) condition, one response key was assigned to unpleasant words and flower stimuli, and the other to pleasant words and insect stimuli. The researchers found that reaction times were shorter for associated stimuli (i.e., the consistent condition) than for nonassociated stimuli (i.e., the inconsistent condition.). This, in essence, is the IAT effect. Typically, the reaction time differential rather than the response accuracy differential across the two conditions is used as a measure of differences in the strength of stimulus associations. Social cognitivists use the existence of these IAT effects to draw inferences about unconscious cognitive activity or hidden prejudices and beliefs.

The IAT has become a favorite instrument of social-cognitive researchers interested in studying attitudes, due to its claims of being resistant to intentional deception by subjects and its ability to measure attitudes of which subjects themselves may not be conscious. However, it has also been met with several criticisms. Most of these criticisms are concerned with the ambiguity of its terminology, the nature of its scoring process (e.g., De Houwer, 2006; Fiedler, Messner, & Bluemke, 2006; Govan, & Williams, 2004; Karpinski, & Hilton, 2001; Olson & Fazio, 2003; Rothermund & Wentura, 2004; Steffens & Plewe, 2001) and its use of statistical inference (see also Blanton & Jaccard, 2006). Others have provided evidence that participants can intentionally influence IAT effects (e.g., De Houwer, Beckers & Moors, 2007; Fiedler & Bluemke, 2005; see also Nosek & Banaji, 2001 for a review).

Unsurprisingly, behavior analysts have also critiqued the IAT's use of mentalisms and poorly understood stimulus control procedures (see Gavin et al., 2008; Gavin et al., in press; Roche et al., 2005; Ridgeway, et al., 2010). A behavioral model of the IAT was offered by Roche et al. (2005) and was later tested empirically by Gavin et al. (2008) and Ridgeway et al. (2010). According to this model, the IAT works by measuring the ease with which common response functions can be established for exemplars from different verbal categories (e.g., *African American* and *Good*). It also compares the accuracy and speed of responding in the acquisition of these common functions with those recorded for the establishment of incompatible response functions for members of the same verbal categories (e.g., *European* and *Good*; see also O'Toole & Barnes-Holmes, 2007).

The behavioral model of the IAT is critical of several features of the IAT presentation format, and it is important that we briefly outline some of these concerns because these are all addressed and improved upon in the FAST format outlined in the current study. First, in the experimental analysis of behavior, reaction times alone are not typically treated as indicative of the strength or stability of any instance of behavior. Instead, response accuracy or fluency (combined accuracy and speed) measures are usually used for this purpose. The FAST, therefore, relies on a fluency-type measure—the number of successive correct responses produced on trials during which a strictly limited response window is enforced. This reliance upon fluency over response time also circumvents several problems regarding the manner in which response times are calculated using the IAT (see Gavin et al., 2012, and Ridgeway et al., 2010, for detailed critiques).

Second, the putative feedback presented during the IAT may in fact function as a form of punishment. Subjects are informed only when erroneous responses are produced, by the presentation of a red X on screen. No feedback is provided following correct responses. This is a less than efficient way to establish response fluency, and the delivery of punishment has unknown effects on subsequent responses and response rates. This

imbalanced feedback procedure makes the emergence of IAT effects more likely, but for reasons not clarified by its inventors. More specifically, this procedure ensures that whichever of the two test blocks contains fewer errors (the relationally inconsistent block) will also be the block that involves less punishment, therefore possibly exaggerating the accuracy and reaction time differences across blocks in the expected direction. This exaggeration may occur because rapid responding in the consistent block is negatively reinforced by the removal of both interruptions to response fluency and the omission or reduction of negative feedback. In contrast, rapid responding in the inconsistent block is punished, thereby leading to slower responding. Previous research has shown that response caution during difficult IAT tasks can partially explain IAT effects (Klauer, Voss, Schmitz, & Teige-Mocigemba, 2007). We would suggest that imbalanced and negative feedback during the IAT produces precisely such imbalanced response caution. Therefore, in the current FAST procedure, feedback follows both correct and incorrect responses.

Third, the most popular IAT scoring technique (D-algorithm; Greenwald, Nosek, & Banaji, 2003) involves a standardized reaction times measure (taken across a fixed trial block) as the core index of association strength (or what we would call stimulus relation strength). However, the trajectories of reaction times across trials are not assessed, and response rates and changes in rate are usually irrelevant to the measures (see also Blanton & Jaccard, 2006). In line with behavior analytic tradition, the FAST format measures learning *rates* to predetermined criteria (i.e., block lengths will vary across subjects) and takes learning rate as the primary behavioral measure over standardized latency scores.

Fourth, the IAT's scoring method involves data cut-off points and subject elimination procedures designed to stabilize data and increase statistical significance in subsequent analyses. Thus, the scoring method is psychometric in style, insofar as psychometrics involves the creation of data stability through means other than improved stimulus control. In contrast, the stability of subjects' behavior during the FAST will be enhanced across research studies using methods designed to improve experimental control over response variability, the analysis of which is an important part of our subject matter (Sidman, 1960).

The current FAST procedure emerged directly from research into developing functional IAT variants within the experimental analysis of behavior literature (e.g., Gavin et al. 2008; Gavin et al., 2012; Ridgeway et al., 2010). However, we should point out that one other behavior analytic alternative to the IAT has also been developed in parallel (i.e., the implicit relational assessment procedure; see Barnes-Holmes, Barnes-Holmes, Stewart, & Boles, 2010). We will consider this test format and its relevance to the current FAST procedure in the Discussion.

The current study involved the examination of the Function Acquisition Speed Test (FAST) procedure for assessing the existence and strength of laboratory controlled stimulus relations. It uses as its core index the rate at which a simple functional class can be established (i.e., the number of training trials required to produce a functional response class at a given level of fluency). More specifically, in the current experiment, subjects will be exposed to a matching-to-sample procedure in which relations between nonsense syllables will be established (i.e., A1–B1, A2–B2; Phase 1). Next, subjects will complete a testing procedure consisting of three phases. Two of these phases will be baseline phases that will measure the number of training trials required to establish simple discrimination response classes when the two discriminative stimuli involved are novel and previously unrelated (i.e., Phases 2 and 4). Phase 3 will consist of two testing blocks. The first block will measure the rate at which two response classes can be established when the discriminative stimuli participating in one of the functional stimulus classes are already previously related as a result of the matching-to-sample training (i.e., A1 and B1 will share common response functions). The second block will

measure the acquisition rate for two response classes that are orthogonal to the trained stimulus–stimulus relations (i.e., A1 and B1 will *not* share common response functions). The learning rates observed for these two blocks will then be compared to each other and the baseline acquisition rates (using entirely novel and unrelated stimuli) to determine if the resulting FAST indices are sensitive to the preexisting stimulus relations.

Method

Subjects

Twenty-one volunteers participated in the study as subjects. Three were eliminated due to their performance during the FAST or the baseline blocks (see Results). Of the 18 remaining volunteers, 6 were male and 12 were female. Ages ranged from 19 to 36 years ($M = 22.94$, $SD = 3.72$). The subjects were recruited both from acquaintances of the experimenter and using a “snowballing” technique whereby subjects were asked to recruit a further subject for participation. Informed consent was obtained in writing from all subjects.

Apparatus

All phases of the experiment were presented to subjects on an Apple MacBook laptop computer with a 13-inch monitor (1024×768 pixel resolution). Stimulus presentations were controlled by the software package Psyscope (Cohen, MacWhinney, Flatt, & Provost, 1993), which also recorded all responses. Stimuli consisted of 14 nonsense syllables (ter, lar, jum, mip, ler, mau, hox, yun, pim, kon, geq, kav, zuv, jin), randomly assigned to their roles as samples, comparisons, and FAST stimuli (see below). These will be referred to in future using alphanumerices.

General Experimental Procedure

The experiment consisted of four phases. Phase 1 (stimulus matching procedure) required approximately 8 minutes to complete. Each of the four subsequent function training blocks (Phases 2–4) typically required 2 to 4 minutes to complete. Phases 2 and 4 each consisted of exposure to a single block of baseline function acquisition training, employing novel nonsense stimuli. Phase 3 involved the presentation of the FAST. This consisted of two blocks; a block testing for the speed of acquisition of a functional response class containing related stimuli (i.e., a “consistent” test block), and a block testing for the speed of acquisition of a functional response class containing stimuli from distinct stimulus relations (i.e., an “inconsistent” test block). There was a maximum of 100 trials presented for all baseline and FAST blocks. Subjects who could not reach criterion within this limit were excluded from the study. Three subjects were eliminated on this basis (see Results).

All phases were presented consecutively by the computer software, which also controlled the delivery of instructions at the beginning of each phase. Subjects sat comfortably at a standard computer desk and viewed the computer screen at a distance of approximately 60 to 70 cm and at eye level.

Phase 1: Matching-to-sample training. In this first phase, subjects were exposed to four matching-to-sample training tasks, each designed to establish two simple stimulus relations. The relations trained were: choose B1, not B2, when A1 is present; choose B2, not B1, when A2 is present, choose A1, not A2, when B1 is present; and choose A2, not A1, when B2 is present (see Figure 1). The purpose of this phase was to establish two laboratory controlled stimulus relations that could be employed to assess the utility of the FAST in determining the existence and strength of stimulus relations established during Phase 1.




Figure 1. The stimulus relations trained during Phase 1. Solid lines indicate reinforced matching responses.

Subjects were presented with the following instructions at the onset of Phase 1:

In a moment, some words will appear on this screen. Your task is to look at the word at the top of the screen and choose one of the two words at the bottom of the screen by “clicking” on it using the computer mouse and cursor. During this stage, the computer will provide you with feedback on your performance. You should try to get as many answers correct as possible. If you have any questions, please ask them now. When you are ready, please click the mouse button.

All trials were presented on the computer screen against a white background. A trial began with the presentation of the sample stimulus at the top center of the screen in black 24-point font. One second later, the two comparison stimuli were displayed in black 24-point font in the bottom left and right corners of the screen (see Figure 1). The positions of the comparison stimuli were counterbalanced across trials. All stimuli remained onscreen until the subject responded by clicking on one of the comparison stimuli using the computer mouse and cursor. Immediately upon a response, the screen cleared and corrective feedback (“Correct” or “Wrong” in red 48-point font) appeared in the center of the screen for 1.5 s. Trials were presented in blocks of 32 (i.e., each of the four trials presented eight times). Subjects were required to complete successive training blocks until a criterion of 31 or 32 correct responses in a single 32-trial block was met (i.e., 96.9% correct).

Phases 2, 3, and 4 comprised the full FAST procedure, which included two baseline phases (Phases 2 and 4) and a test phase (Phase 3), which in turn consisted of both a consistent and inconsistent test block. For clarity, we will describe the Phase 3 test blocks first.

Phase 3: Function Acquisition Speed Test. In this phase, subjects were exposed to the two-block FAST using the A1 and B1 stimuli from Phase 1 and two further novel nonsense syllable stimuli, N1 and N2. One of these blocks (the consistent block) established two functional stimulus classes (A1–B1 and N1–N2) that were consistent with the stimulus–stimulus relations established in Phase 1 (i.e., in which A1 was matched with B1). The other block (inconsistent) established two functional stimulus classes (A1–N1 and B1–N2) that were inconsistent with the relations established in Phase 1. Figure 2, middle panel, illustrates the tasks involved in each block. The order of the consistent and inconsistent blocks was randomized across subjects.

Subjects were presented with the following instructions at the onset of each FAST block:

In the following section, your task is to learn which button to press when a word appears on screen. **IMPORTANT:** During this phase you should press only the A key or the J key. Please locate them on the keyboard now. This part of the experiment will continue until you have learned the task and can respond without error. To help you learn, you will be provided with feedback telling you if you are right or wrong. If you have any questions, please ask the researcher now. Press any key when you are ready to begin.




Figure 2. A schematic of the various response functions assigned to stimuli during baseline and FAST blocks. The three panels, read from top to bottom, also represent the sequence of presentation of the two baselines blocks and the two FAST blocks.

All trials were presented on the computer screen with a white background. A trial began with the presentation of one of four nonsense syllable stimuli (A1, B1, N1, N2) in the center of the screen in 48-point black font. The stimuli remained on-screen for a period of 3 s or until a response was emitted (i.e., a 3 s response window was enforced). Each of the four stimuli was presented in a quasirandom order in blocks of four trials (i.e., no more than two consecutive exposures to any one stimulus was possible).

Immediately upon the production of a response, corrective feedback was presented (i.e. either “Correct” or “Wrong” presented in red 48-point font in the center of the screen for 1.5 s). If no response was emitted within the 3 s response window, an incorrect response was recorded, but no feedback was provided. In that case, the subsequent trial began immediately upon the end of the 3 s response window. Subjects were exposed to trials until 10 consecutive correct responses were produced. A predetermined limit of 100 trials was enforced because pilot research had indicated that once this limit was reached, the subject was unlikely to complete the block before giving up or being asked to cease by the experimenter.

In summary, the FAST blocks attempted to establish two response classes under two conditions; one in which previously related stimuli participated in the same functional stimulus class, and one under which they participated in distinct functional stimulus classes.

Phases 2 and 4: Baseline function acquisition rates. This purpose of these phases was to establish a baseline level of response class acquisition using novel and unrelated stimuli against which acquisition rates with target stimuli could be compared. The procedure used during these two phases was identical to that employed during the critical FAST blocks (Phase 3). However, only one block was provided in each phase. In addition, different response keys (Z and M) were employed during baseline phases in order to prevent any conflicting response histories across baseline and critical FAST phases.

Two baseline phases were presented (one before and one following the FAST; Phase 2 and 4, respectively) in order to assess the stability of baseline rates of function acquisition across time. Administering two baseline phases also had the advantage that it would allow

for the calculation of a mean baseline acquisition rate if baseline performances proved to be unstable across time (see Results). Figure 2 (top and bottom panel) illustrates the tasks involved. The following instructions were delivered at the start of each baseline phase:

In the following section, your task is to learn which button to press when a word appears on screen. **IMPORTANT:** During this phase, you should press only the Z key or the M key. Please locate them on the keyboard now. This part of the experiment will continue until you have learned the task and can respond without error. To help you learn, you will be provided with feedback telling you if you are right or wrong. If you have any questions, please ask the researcher now. Press any key when you are ready to begin.

The two baseline blocks involved novel and unique stimuli. Phase 2 employed X1, X2, Y1, Y2 as stimuli while Phase 4 employed X3, X4, Y3, Y4 as stimuli. These phases attempted to establish the following relations, respectively: X1–Y1, X2–Y2 and X3–Y3, X4–Y4.

Results

Of the 21 volunteers who completed the experimental procedure, the performances of three subjects were not included in the data analysis. Specifically, Subjects 2 and 6 were eliminated on the basis that they failed to reach response fluency on one of the baseline blocks within 100 trials. Subject 8 failed to reach the fluency criterion on one of the FAST blocks (the consistent block).

Matching-to-Sample Training

All subjects successfully completed matching-to-sample training. The average number of training trials needed to reach the criterion (97% correct) was 65.61. Table 1 illustrates the training requirements for each subject.

Table 1
Number of Trials Required to Complete Matching-to-Sample Training (Phase 1)

Subject	1	3	4	5	7	9	10	11	12	13	14	15	16	17	18	19	20	21
No. of Trials	64	32	32	32	128	32	96	64	64	64	64	64	64	64	128	32	64	96

Baseline Blocks

Table 2 shows the number of trials required by each subject to reach the fluency criterion (10 correct responses in a row without error) for each baseline block (Phases 2 and 4). These data illustrate what appears to be large variation in performance from the first to the second baseline block. Upon close inspection, however, this variation would not appear to be consistent, and cannot therefore be considered a practice effect. A Wilcoxon signed ranks tests found the difference in trial requirements across the two baseline blocks to be nonsignificant ($z = -1.198, p = .231$, two-tailed). Given the within-subject variability in performance across the baseline blocks, the researchers calculated the mean number of trials to criterion across both baseline blocks for each subject. They then used this figure as the baseline acquisition rate against which to compare FAST block acquisition rates.

Table 2
Number of Baseline Trials Required to Reach the Fluency Criterion Across Both Baselines (Phases 2 & 4)

Subject	1	3	4	5	7	9	10	11	12	13	14	15	16	17	18	19	20	21
Baseline 1	26	84	12	41	40	27	42	70	24	88	20	48	66	21	35	49	20	12
Baseline 2	31	26	10	15	48	12	12	12	100	31	58	30	16	12	24	58	59	21
Mean	28.5	55	11	28	55	19.5	27	41	62	59.5	39	39	41	16.5	29.5	53.5	39.5	16.5

Function Acquisition Speed Test

Each subject completed two FAST blocks (Phase 3), a consistent and an inconsistent block. Table 3 shows the number of trials to criterion for each test block and the order in which the FAST blocks were administered. It was expected that a larger number of acquisition trials would be required for subjects to reach criterion on the inconsistent block compared to the consistent block. A Mann–Whitney *U* test indicated that there were no significant order effects observed in the data. That is, trial requirement differentials across the FAST blocks were not significantly different *based on order* ($z = -0.442, p = .67$, two-tailed).

Of the 18 subjects, 13 showed a faster rate of response function acquisition in the consistent block compared to the inconsistent block, as expected. Four subjects showed a small acquisition rate differential in the unexpected direction, while one subject (Subject 7) showed no acquisition differential across the FAST blocks. The mean differential in FAST block trial requirements in the expected direction was 27.1. In contrast, the mean differential in FAST block trial requirements in the unexpected direction was considerably smaller at -7.5, indicative of the expected FAST effect at the group level.

Table 3
Number of Trials Required to Reach the Fluency Criterion Across Each of Two Blocks of the FAST (Phase 3)

Subject	1	3	4	5	7	9	10	11	12	13	14	15	16	17	18	19	20	21
Inconsistent Block	88	60	40	20	22	25	32	15	43	40	10	76	23	47	32	100	22	22
Consistent Block	15	41	10	16	22	10	24	12	48	28	12	26	30	10	20	17	16	38
Difference	73	19	30	4	0	15	8	3	-5	12	-2	50	-7	37	12	83	6	-16
Order	1	2	1	2	1	2	2	2	2	2	1	1	2	1	1	2	1	1

Note. The order in which the two test blocks were administered is also indicated (1 indicates that the inconsistent block was administered first, whereas 2 indicates that the consistent block was administered first).

A Wilcoxon signed ranks tests was performed on the acquisition rate data during the two FAST blocks. The acquisition rate differential across the FAST blocks was found to be significant ($z = -2.580, p = .005$, one-tailed), indicating a strong FAST effect at the group level.

In order to quantify the magnitude of each subject’s own fluency differential in response class acquisition across the FAST blocks, we employed a simple Strength of Relation (SoR) Index. This was calculated for each subject by subtracting the number of trials required to complete the consistent block of the FAST from the total number of trials required to complete the inconsistent block of the FAST, and dividing this differential by the mean number of trials required by that subject to reach criterion during the baseline blocks. In effect, the SoR index can be used here as a measure of the strength of the A1–B1 stimulus matching response, insofar as each subject’s acquisition rate differentials across FAST blocks is adjusted by a factor representing their own unique baseline acquisition rates for such task types. Larger positive SoR indices indicate higher strength stimulus relations (i.e., A1–B1 stimulus matching responses are fast to acquire and are more resistant to disruption by competing reinforcement contingencies), while zero or negative responses indicate absent or reversed FAST effects (i.e., A1–B1 relations are slow to acquire and are more easily disrupted by competing reinforcement contingencies).

Table 4
Strength of Relation [SoR] Indices

Subject	1	3	4	5	7	9	10	11	12	13	14	15	16	17	18	19	20	21
[SoR]	2.56	0.35	2.73	0.14	0	0.77	0.3	0.07	-0.08	0.2	-0.01	1.28	-0.17	2.24	0.41	1.55	0.152	0.97

Note. Larger numbers indicate stronger A1–B1 stimulus matching responses (i.e., stimulus relations). Negative numbers indicate a reversal of the expected acquisition rate differential across FAST blocks (i.e., a reversed FAST effect).

Table 4 shows that for the majority of subjects, the FAST produced positive identifications of a history of A1–B1 training (i.e., it was sensitive to the subjects' histories of relating A1 and B1). That is, positive SoR indices were calculated for 13 of the 18 subjects. This table also shows that where SoR indices were negative, indicating a reversed acquisition rate differential across the FAST blocks, the magnitude of the index was small. In other words, the tendency is predominantly toward large positive SoR indices and small negative ones.

A one-sample Wilcoxon was also conducted to assess the significance of SoR index scores against zero. This analysis showed that SoR index scores differed significantly from zero in the expected positive direction ($z = 2.580, p = .005$, one-tailed), indicating a history of A1–B1 relational training for the subject group as a whole and a significant FAST effect.

The researchers then examined whether or not poorly established A–B relations during Phase 1 may partially account for those instances in which the predicted FAST effect was not observed. Some subjects who failed to show a FAST effect had required a large number of training trials (128, Subject 7, SoR = 0), while others who failed to show a FAST effect had required a smaller number (64, Subjects 12, 14, and 16). However, none of those who failed to show a FAST effect had completed the training in the minimal number of trials (32). It was suspected that requiring a larger number of training trials to complete simple conditional discrimination training may be indicative of poor stimulus control. In simple terms, the more training subjects required to pass Phase 1, the more likely it was that fortuitous forms of stimulus control (or mere chance-level choices during matching-to-sample tasks) controlled final performances. To check this idea, inferential correlational analyses were conducted to examine the relationship between the number of training trials required in Phase 1 and the final FAST effect. A Spearman's rho analysis indicated that training trial number was moderately negatively correlated with the raw trial requirement differential across the FAST blocks ($r = -0.495, n = 18, p = .018$, one-tailed). Training trial number was also found to be moderately negatively correlated with the SoR index of the FAST effect ($r = -0.439, n = 18, p = .034$, one-tailed). Thus, more rapid acquisition of the baseline relations during Phase 1 was moderately associated with a larger FAST effect.

In summary, 13 out of 18 of the subjects showed the expected FAST effect. At a group level, subjects required significantly fewer training trials for consistent response classes to be established compared to inconsistent response classes. The SoR index also provided a statistically significant measure of preexisting stimulus–stimulus relations at the group level. Finally, rapid acquisition of baseline relations during Phase 1 was found to be moderately associated with larger FAST effects.

Discussion

The current study established that a FAST can be used to determine the preexistence of specific stimulus–stimulus relations at the group level. This result is broadly comparable with the types of group-level effects observed with other implicit tests. Thus, the current procedure might form the basis of a novel behavioral methodology for assessing the existence, and possibly the strength, of previously existing stimulus relations in the vernacular. In this regard, the FAST technique may prove useful to researchers wishing to assess stimulus relations using a simple and functionally transparent implicit testing-style methodology.

The FAST procedure may also be of use in an educational or research context in which the strength of a relational response may be of interest. At present, behavior analysts primarily utilize the transition from inaccurate to accurate responding (i.e., percentage correct) as the measure of matching response strength. However, several studies have used response times to measure relational responses at the level of conditional discrimination and derived relations (e.g., Bentall, Dickins, & Fox, 1993; O'Hara, Roche & Barnes-Holmes, 2002; Roche, Linhehan, Ward, Dymond, & Rehfeldt, 2004; Steele & Hayes, 1991; Wulfert & Hayes, 1988) because speed of responding has been shown to be a valid index

for differentiating response strengths when response accuracy has stabilized (Johnson & Layng, 1992; see also Spencer & Chase, 1996). The FAST method allows for the assessment of a matching response in contexts, in which, using traditional accuracy or response time criteria alone, the matching response is either already at maximum strength or not at sufficient strength to be considered established. In effect, the procedure would allow for the concurrent assessment of a matching response repertoire, even as the very repertoire is emerging or even after it has reached maximum accuracy. Measuring the ongoing emergence of relations in terms of relation strength may be useful when individuals display slow learning curves and where even slight improvements in relational responding are of interest. In the context of implicit testing, the FAST method would allow for an analysis of verbal relations between stimuli that have not yet emerged at sufficient strength in the repertoire of the subject (i.e., emerging racist forms of speech) to be discriminated by either the subject or an observer, or recorded reliably using accuracy-based measures.

The FAST procedure lived up to its name in the current experiment by generating the predicted effect within a matter of a few minutes. While the number of trials delivered to each subject is variable by design, a typical FAST, including the administration of two baseline blocks, required approximately 5 minutes to complete (this time does not include the time taken to establish the target A–B stimulus relations, a procedure which is unnecessary when using the FAST to assess preexisting, real-world stimulus relations). This is considerably less time than is required to complete another behavior analytic test; the Implicit Relational Assessment Procedure (IRAP; Barnes-Holmes, Barnes-Holmes et al., 2010). The IRAP is an exhaustive test for the varieties of stimulus relations that obtain between exemplars of a complex verbal stimulus relation. In contrast, the FAST is intended as a rapid and easily administered “litmus test” for the existence and possibly strength of a specific stimulus relation between two verbal stimuli. It is important to point out, therefore, that there is a trade-off to be made between speed of administration and depth of analysis when choosing between a FAST and an IRAP procedure.

The variation in learning rates across the two baseline exposures and across individual subject’s FAST effects is a cause for concern. This instability in the baseline block data is not due to practice effects, which would have been evidenced by systematic decreases in trial requirements from Baseline 1 to Baseline 2. Further research is required to identify the source of this variation, which may have partly contributed to the reversal of the predicted FAST effect for four of the subjects. Of course, one might suggest that the most parsimonious explanation for the observed behavioral variation is that the FAST effect is superimposed by natural “noise.” However, this sort of explanation will not suffice for behavior analysts interested in identifying a source of behavioral variability for its own sake.

Response stability might be improved by several exposures to individual baseline blocks at the outset of a testing session. These blocks could function not only as additional baseline measures but as practice blocks leading to behavioral stability before the crucial FAST blocks are administered. It is important to remember, however, that despite this apparently random variability, predicted effects still emerged across the two key FAST blocks in the current study. Thus, any reduced variability in FAST block performance achieved in future research would likely serve only to enhance the FAST effect rather than diminish it.

Reduced response windows may also aid in increasing behavioral stability. This idea can easily be tested empirically. However, the current data and other published studies suggest that such a modification to the FAST may be fruitful. First, response latency data in the current study suggested that most responses occurred within 2,000 ms. In effect, the response window of 3,000 ms was not a pressing constraint on response speed. Second, while not directly comparable, other studies have successfully employed short response windows in IAT preparations (e.g., 200 ms; Verschuere, Prati, & De Houwer, 2009) to restrict conscious response strategies. In one priming study in which single stimuli were

presented individually, as in the FAST, a response window of only 600 ms was employed (Degner, 2009).

One means by which the procedural implicitness of the FAST may be enhanced in real-world contexts is by increasing the number of exemplars employed to assess the relations between verbal categories. For instance, several exemplars of positive words might be used in the place of the target category label "fun" to assess the relation between *mathematics* and *fun*. These exemplars would be employed interchangeably across trials in the FAST blocks. Irrelevant stimuli might also be employed as distractors to reduce the probability of subjects identifying the relations of interest. It is important to understand, however, that this would serve only to improve procedural implicitness. Such measures cannot help us to determine whether or not a subject was *previously* aware of the relations under analysis (i.e., outcome implicitness). In fact, in the current study, we can assume that subjects were aware of the A1–B1 relations under analysis, because this was trained and tested explicitly during the conditional discrimination training phase. Further research is required in order to determine if untrained or derived (unconscious) relations between experimental stimuli can be detected using the FAST procedure. This would determine whether or not the FAST procedure achieves both procedural and outcome implicitness.

There is a potential confound present in the current study with regard to what are known as "salience asymmetries" (Rothermund & Wentura, 2004) across the test stimuli. Specifically, the conditional discrimination training may have produced differences in subjects' levels of familiarity with the A and B stimuli compared to the novel nonsense syllables employed as N1 and N2 control stimuli. It is possible that subjects responded to a fortuitous functional equivalence between the A and B stimuli on the grounds that they shared a history of simply being recognizable. Indeed, familiarity of stimuli has been shown to accelerate the acquisition of stimulus equivalence classes (Holth & Arntzen, 1998). In contrast, the novel nonsense syllables denoted as N1 and N2 share their relative novelty. These asymmetries alone might help to account for the observed effects, although we see this as unlikely. Further research should involve the use of control blocks to eliminate the problem of salience asymmetry. For instance, in one control block, A1 and B2 (rather than A1 and B1) might share a response function, while in the other block, A1 and N2 would share a response function. We would predict that acquisition of a common response function for A1 and B2 would be significantly slower than for A1 and N1. This outcome would indicate that A1 and B2 were nonequivalent, despite being equally familiar to subjects. A more thorough examination of the issue would involve engineering equal levels of familiarity across all experimental stimuli prior to the critical testing phases. It is important to remember, however, that in the current context, even if stimulus salience alone were to be the only feature controlling the A1–B1 relation, this would not take from the fact the FAST procedure was sensitive to this relation. Unknown salience asymmetries only pose a problem when researchers make claims regarding the preexperimental associations existing between real-world stimuli. Nevertheless, if the FAST procedure is to be employed in real-world settings, its susceptibility to salience asymmetries must be assessed in further research.

One important challenge for any behavioral test is for it to allow for the discrimination of effects at the level of the individual. In other words, a test such as the FAST should be able to identify a history of stimulus relations for each individual subject. While group-level effects were strong in the current study, the FAST failed to identify the stimulus relation histories of four individuals, who, in fact, showed the reverse effect. It was promising that these reversed effects were among the smallest in the data set, strongly suggesting that they were based on uncontrolled, low-level variability rather than on pure randomness (i.e., no control). Nevertheless, nothing short of 100% predictive validity must be sought in future research if the processes analyzed in the current study are to be fully understood. Achieving that will be no easy task, especially when one considers that such levels of predictive validity have not yet been achieved by even the most mature implicit tests. Nevertheless, given that the analysis of stimulus control and behavioral variability

are the stock in trade of the behavior analyst, we are well positioned to produce a test with near perfect predictive validity, based upon the behavioral processes explicated in the current study. It should be remembered, however, that while the authors do not view the FAST as a psychometric test (i.e., measuring an internal trait), the test will nevertheless display psychometric properties that should be explicated across research studies so that informative comparisons can be made between it and more commonly used tests such as the IAT (see Fiedler et al., 2006, for reliability and validity criteria an implicit test should meet).

The issue of comparing the FAST to other implicit test procedures is relevant to the matter of the validity of the SoR index, employed here to quantify strength of relations between test stimuli. Specifically, it could be suggested that in the wider field of implicit testing, such a scoring method, may not satisfy popular normative assumptions. That is, raw data from blocks of implicit tests are usually normalized in some way, although this very fact has been met with criticism from behavior analysts (Gavin et al., 2008, 2012; Ridgeway et al., 2010) and social cognitivists (Blanton & Jaccard, 2006). The most common method is to calculate z scores, following trimming of the raw data set to remove outlying data points and outlying overall subjects (this method is common to both the IAT and IRAP). This procedure will of course stabilize data and increase p values and effect sizes using Cohen's (1988) criteria. Behavior analysts are normally wary of such statistical treatments because the purpose of data-cleaning procedures is to create data stability post hoc rather than through improved stimulus control measures. It is for this reason that the SoR index employed here does not involve data normalization, in the usual sense, but nevertheless corrects raw acquisition rate differences across test blocks by a factor representing the individual subject's fluency at that particular task type (i.e., the untreated baseline acquisition rate). Thus, while this technique does not constitute data normalization in the usual sense, it is important to understand that it was not intended to do. Nevertheless, the fact still remains that some researchers may wish to compare the FAST effects with those derived from other preparations, such as the IAT or IRAP, both of which employ normalized difference score (i.e., across test blocks) algorithms to quantify test effects (e.g., Greenwald et al., 2003). It is difficult to say at present how FAST effects might compare to those derived using other procedures and scoring systems.

An important asset of the FAST procedure is that it provides a nonrelative measure of stimulus relations. That is, it allows us to test the existence and possibly the strength of a relation between a target word and any other word, without needing to also measure relations to control words for comparison purposes (e.g., an IAT might assess that for a group of subjects *Black* goes with *Bad* to a greater extent than *White* goes with *Bad*). In this regard, it is a preferable technique to the IAT. While nonrelative measures are also provided by the Go/No-Go association task (GNAT: Nosek & Banaji, 2001) and the IRAP (Barnes-Holmes, Barnes-Holmes, et al., 2010), the current method also has the advantage in that it can provide a single index of relation strength that can be used for comparative purposes across research studies.

Because the FAST assesses the association between only two stimuli, it is more specific than most other implicit measures. Furthermore, because there is only one relevant association under analysis, no assumptions need to be made regarding how a number of specific associations are weighted against each other. Such assumptions have proved problematic for the IAT, in which up to four associations between two concepts and two attribute categories are assumed to contribute to the observed effect (Blanton, Jaccard, Gonzales, & Christie, 2006).

The high degree of stimulus-specificity in the FAST raises the possibility that, in real-world settings, using real words as stimuli, test effects may rest to some extent upon the representativeness of the stimuli chosen for the target categories of interest. For instance, an IAT designed to assess relations between African American names and negative words will produce slightly different results depending on the stimuli chosen for this purpose. This problem is known as the "stimulus specificity problem" (see De Houwer, 2001, for a

distinction between a relevant and an irrelevant feature account). The IAT employs multiple exemplars of the stimulus categories of interest (e.g., typically, approximately eight exemplars from each category are employed interchangeably across trials), but the key category labels (e.g., the term *African American*) cannot be used as target stimuli because they are used in the response rules presented concurrently onscreen. Thus, the IAT suffers from the problem of stimulus specificity, while at the same time partially controlling for it by using multiple exemplars to represent the stimulus categories of interest. The FAST, however, is suited to the use of category labels as target stimuli, a strategy suggested by Steffens, Kirschbaum, and Glados (2008) to deal with the very problem of stimulus specificity. However, although, using category labels as target stimuli increases the specificity of the test and conclusions derived from results, it might compromise procedural implicitness by potentially alerting subjects to the categories of interest.

The FAST procedure improves upon the implicitness of other popular measures, through its use of a nonassociative methodology. The use of this methodology also challenges popular associative accounts of implicit test effects popular in the cognitive and social-cognitive literature. Before we outline these potentially important issues, we will explain what we mean by an *associative methodology*.

Popular implicit tests almost universally depend on the contiguous presentation of two stimuli of interest on screen in the form of a rule (e.g., the IAT) or in sequence (e.g., affective priming tasks). This association is manipulated across trial types in order to assess the existence or strength of the preexisting association (or relation) between these two stimuli. In the IAT, for instance, the rule "Press left for *Bad* and *African American*" during one block of testing may be changed to "Press left for *Good* and *African American*" in a subsequent block. The manipulation of this association (e.g., the simultaneous presentation of *Bad* or *Good* and *African American*) is the basis of the response-time differential observed across task blocks. In effect, an association is created between pairs of stimuli in different configurations to assess their compatibility with assumed preexisting "mental associations" (see De Houwer, 2006).

As another example, the affective priming technique involves the presentation of target stimuli immediately following affective prime stimuli (Fazio, Sanbonmatsu, Powell, & Kardes, 1986). Subjects are required to rapidly classify these target stimuli. These classification responses are in turn influenced by the nature of the prime. For instance, if the prime and target are affectively compatible in the subject's history (e.g., *Bad* and *African American*), then reduced response times are observed for the identification of the target (i.e., greater response fluency). If, however, the prime and target are affectively incompatible (e.g., *Good* and *African American*), response latencies are elongated. As with the IAT, stimulus pairs in affective priming tasks are presented in an associative manner, and implicit knowledge is inferred from response differentials across different association configurations. Several other common procedures, such as the emotional Stroop (see Williams et al., 1988), the Go/No-Go association test (Nosek & Banaji, 2001) and the Extrinsic Affective Simon Test (EAST; De Houwer, 2003) employ comparable associative methodologies.

The main disadvantage of associative methodologies in popular implicit tests is that such a procedure provides an opportunity for subjects to discriminate the very relations under analysis (e.g., that between *African American* and *Good* or *Bad*). Behavior can therefore either become subject to social desirability, or subjects may even succeed at a response-faking attempt (e.g., purposely respond in error or more slowly during one trial block). In simple terms, if two category labels, or exemplars of categories of interest, are presented contiguously to a subject, it is likely that many subjects will become aware of the verbal relations or associations of interest to the researcher. This compromises the indirectness or procedural implicitness of the methodology.

IAT research has highlighted the need for improved procedurally implicit measures. Specifically, while it was formerly thought that the IAT was relatively impervious to faking attempts by the subject (Kim, 2003), it now appears that faking attempts by subjects

are more successful on subsequent exposures to the test (Steffens, 2004) and when novel stimuli rather than familiar words are employed (De Houwer, Beckers, & Moors, 2007). Indeed, the susceptibility of the IAT to faking is now generally higher than previously suspected (see Greenwald, Poehlman, Uhlmann, & Banaji, 2009), and several studies have examined the compromises suffered by the IAT as a result of conscious faking efforts (e.g., Fiedler & Bluemke, 2005). Given this, any test feature that reduces the opportunity to discriminate the relations under analysis (i.e., improves procedural implicitness) is a welcome contribution to the literature.

The FAST procedure achieves increased procedural implicitness by presenting only single stimuli in the absence of rules or any instructions involving the target stimuli. Trial-by-trial response reinforcement alone controls responding to stimuli, and only the reinforcement contingencies are manipulated across test blocks. Thus, by employing a nonassociative methodology, the FAST procedure greatly reduces the opportunity for subjects to discriminate the purpose of the test, thereby enhancing its procedural implicitness. Interestingly, recent research has found that the use of a response window, as employed in the current FAST procedure, may also protect against faking attempts by subjects (Verschuere et al., 2009). More specifically, even if a subject becomes aware of the core process of the test, the subject cannot easily alter response time and accuracy when under a strict time constraint. Parenthetically, this recent realization among social cognitivists that enforced response fluency should replace or complement instructions to subjects to “go fast” parallels repeated calls for such an emphasis in the behavior analytic literature (see Gavin, Roche, & Ruiz, 2008; Gavin, Roche, Ruiz, Hogan, & O’Reilly, 2012; Ridgeway et al., 2010; Roche et al., 2005; see also Barnes-Holmes, Murphy, et al., 2010).

The recently developed Implicit Relational Assessment Procedure also has nonassociative stimulus presentation features, insofar as subjects are presented with complex contextually controlled relations rather than simple stimulus pairings (Hughes, Barnes-Holmes, & De Houwer, 2011). For instance, a racial bias IRAP might involve presenting the words *Black* and *Bad* on a computer screen simultaneously and require subjects to respond to these stimuli as relationally compatible or incompatible, as instructed. While the contiguous relation between the target stimuli is not manipulated, as it typically is in associative methodologies, the response requirements (i.e., the instructions) are. This feature has implications for the nature of the associations being assessed by implicit tests of that kind. That is, it challenges the popular assumption that implicit tests measure only simple, direct associations in memory rather than syllogisms, indirect knowledge or propositional knowledge (see Hughes et al., 2011). Propositional knowledge is usually often assumed to be assessed more readily using explicit measures, such as paper-and-pencil tests. Nevertheless, the IRAP’s nonassociative procedure does not improve upon the procedural implicitness of more popular implicit tests. That is, because of the simultaneous presentation of target stimuli on screen and in instructions, subjects may still easily discriminate the relations under analysis in the test.

The foregoing also raises the interesting conceptual issues regarding the types of stimulus associations being assessed by the FAST. Because the FAST procedure does not present stimulus pairs contiguously or contingently, test effects challenge the usual explanatory mechanism provided in social-cognitive terms. That is, most implicit test researchers assume that associative methodologies are required in implicit tests because it is precisely these types of associations that one would wish to activate in memory (see Mitchell, De Houwer, & Lovibond, 2009). Test features, such as instructed time pressure, ensure that spontaneous and automatic (nonconscious) responding occurs that is controlled by these primitive mental associations. However, the FAST procedure cannot be easily viewed as activating mental associations in the usual way. While the current study did involve the creation of simple stimulus associations during matching-to-sample training, stimulus pairs were never presented in the FAST in any contiguous or contingent manner, and there were no instructions presented at any stage in which the target stimuli were

referred to. Of course, an interpretive mediational cognitive account can always be easily constructed. One possibility, for instance, is that the conditional discrimination training established associations between A1 and B1 that lead to each evoking mental representations of the other during the FAST. Thus, A1 and B1 may each have become a cue for performing the response required in the presence of the other. In the consistent FAST block, this process may have facilitated the function training, whereas in the inconsistent block it may have militated against it, thereby explaining the overall test effect (see Hall et al., 2003, for precisely such an account). Of course, while acceptable to those of a cognitive persuasion, this account is entirely speculative and is not fitting with the behavior analytic perspective adopted by the current research (see also Hughes et al., 2011; Smyth, Barnes-Holmes, & Barnes-Holmes, 2008).

Other common interpretive explanations applied to the IAT, and potentially applicable to the FAST, appeal to sources of control that are more amenable to functional analysis. For instance, IAT effects are sometimes explained in terms of S-R incongruity (e.g., the same response for negatively evaluated *mathematics* and positively evaluated *fun*; De Houwer, 2001), or the trial-to-trial inertia effect of stimulus feature-response associations, their inhibition in case of interference, and the subsequent carryover of inhibition (Klauer & Mierke, 2005). These are all theoretical views informed by a fundamentally cognitive perspective but which may point to important stimulus control issues amenable to study in the behavioral laboratory. These various cognitive accounts notwithstanding, the important point from our perspective is that the FAST effects observed in the current study are also easily amenable to a behavioral analysis in terms of competing contingencies over relational responding, as offered in the original behavioral model of the IAT (see Gavin et al., 2008). It is important to note that no recourse to mental associations or similar cognitive accounts seems *required* (even if preferred) to explain the current data. In effect, the reliance of the FAST on well-understood behavioral principles, and its employment of simple scoring techniques with face validity (i.e., the SoR index), circumvent the *need* to appeal to hypothetical constructs and processes in explaining the current effects. Thus, rather than concern ourselves with matters of what types of assumed and hypothetical mental associations (associative or propositional) actually produced the current effects, we will focus in subsequent studies on the controlling variables that produced these test outcomes and thereby circumvent such debate altogether.

References

- BANSE, R., SEISE, J., & ZERBES, N. (2001). Implicit attitudes toward homosexuality: Reliability, validity, and controllability of the IAT. *Experimental Psychology, 48*, 1145–1160.
- BARNES-HOLMES, D., BARNES-HOLMES, Y., STEWART, I., & BOLES, S. (2010). A sketch of the Implicit Relational Assessment Procedure (IRAP) and the relational elaboration and coherence (REC) model. *The Psychological Record, 60*, 527–542.
- BARNES-HOLMES, D., MURPHY, A., BARNES-HOLMES, Y., & STEWART, I. (2010). The Implicit Relational Assessment Procedure (IRAP): Exploring the impact of private versus public contexts and the response latency criterion on pro-white and anti-black stereotyping among white Irish individuals. *The Psychological Record, 60*, 57–66.
- BENTALL, R. P., DICKINS, D. W., & FOX, S. R. A. (1993). Naming and equivalence: Response latencies for emergent relations. *The Quarterly Journal of Experimental Psychology, 46B*, 187–214.
- BLANTON, H., & JACCARD, J. (2006). Arbitrary metrics in psychology. *American Psychologist, 61*, 27–41.
- BLANTON, H., JACCARD, J., GONZALES, P., & CHRISTIE, C. (2006). Decoding the Implicit Association Test: Implications for criterion prediction. *Journal of Experimental Social Psychology, 42*, 192–212.

- COHEN J. D., MACWHINNEY B., FLATT M., & PROVOST J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments, and Computers*, 25, 257–271.
- DEGNER, J. (2009). On the (un)controllability of affective priming: Strategic manipulation is feasible but can possibly be prevented. *Cognition & Emotion*, 23, 327–354.
- DE HOUWER, J. (2001). A structural and process analysis of the Implicit Association Test. *Journal of Experimental Social Psychology*, 37, 443–451.
- DE HOUWER, J. (2003). The Extrinsic Affective Simon task. *Experimental Psychology*, 50, 77–85.
- DE HOUWER, J. (2006). What are implicit measures and why are we using them. In R. W. Wiers & A. W. Stacy (Eds.), *The handbook of implicit cognition and addiction* (pp. 11–28). Thousand Oaks, CA: Sage.
- DE HOUWER, J., BECKERS, T., & MOORS, A. (2007). Novel attitudes can be faked on the implicit association test. *Journal of Experimental Social Psychology*, 43, 972–978.
- DIXON, M., R., REHFELDT, R. A., ZLOMKE, K. R., & ROBINSON, A. (2006). Exploring the development and dismantling of equivalence classes involving terrorist stimuli. *The Psychological Record*, 56, 83–103.
- FAZIO, R. H., SANBONMATSU, D. M., POWELL, M. C., & KARDES, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, 50, 229–238.
- FIEDLER, K., & BLUEMKE, M. (2005). Faking the IAT: Aided and unaided response control on the implicit association tests. *Basic and Applied Social Psychology*, 27, 307–316.
- FIEDLER, K., MESSNER, C., & BLUEMKE, M. (2006). Unresolved problems with the “I,” the “A” and the “T”: Logical and psychometric critique of the Implicit Association Test (IAT). *European Review of Social Psychology*, 17, 74–147.
- GAVIN, A., ROCHE, B., & RUIZ, M. R. (2008). Competing contingencies over derived relational responding: A behavioral model of the implicit association test. *The Psychological Record*, 58, 427–441.
- GAVIN, A., ROCHE, B., RUIZ, M. R., HOGAN, M., & O'REILLY. (2012). A behavior-analytically modified implicit association test for measuring the sexual categorization of children. *The Psychological Record*, 62, 55–68.
- GOVAN, C. L., & WILLIAMS, K. D. (2004). Changing the affective valence of the stimulus items influences the IAT by redefining the category labels. *Journal of Experimental Social Psychology*, 40, 357–365.
- GREENWALD, A. G., MCGHEE, D. E., & SCHWARTZ, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74, 1464–1480.
- GREENWALD, A. G., NOSEK, B. A., & BANAJI, M. R. (2003). Understanding and using the implicit association test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85, 197–216.
- GREENWALD, A. G., NOSEK, B. A., BANAJI, M. R., & KLAUER, K. C. (2005). Validity of the salience asymmetry interpretation of the IAT: Comment on Rothermund and Wentura (2004). *Journal of Experimental Psychology: General*, 134, 420–425.
- GREENWALD, A. G., POEHLMAN, T. A., UHLMANN, E., & BANAJI, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, 97, 17–41.
- HALL, G., MITCHELL, C., GRAHAM, S., & LAVIS, Y. (2003). Acquired equivalence and distinctiveness in human discrimination learning: Evidence for associative mediation. *Journal of Experimental Psychology: General*, 132, 266–276.

- HOLTH, P., & ARNTZEN, E. (1998). Stimulus familiarity and the delayed emergence of stimulus equivalence or consistent nonequivalence. *The Psychological Record*, *48*, 81–110.
- HUGHES, S., BARNES-HOLMES, D., & DE HOUWER, J. (2011). The dominance of associative theorising in implicit attitude research: Propositional and behavioral alternatives. *The Psychological Record*, *61*, 465–498.
- JOHNSON, K. R., & LAYNG, T. V. J. (1992). Breaking the structuralist barrier: Literacy and numeracy with fluency. *American Psychologist*, *47*, 1475–1490.
- KARPINSKI, A., & HILTON, J. L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology*, *81*, 774–778.
- KIM, D. (2003). Voluntary controllability of the Implicit Association Test (IAT). *Social Psychology Quarterly*, *66*, 83–96.
- KLAUER, K. C., & MIERKE, J. (2005). Task-set inertia, attitude accessibility, and compatibility-order effects: New evidence for a task-set switching account of the Implicit Association Test effect. *Personality and Social Psychology Bulletin*, *31*, 208–217.
- KLAUER, K. C., VOSS, A., SCHMITZ, F., & TEIGE-MOCIGEMBA, S. (2007). Process components of the Implicit Association Test: A diffusion-model analysis. *Journal of Personality and Social Psychology*, *93*, 353–368.
- KOHLBERG, B. K., HAYES, S. C., & HAYES, L. J. (1991). The transfer of contextual control over equivalence classes through equivalence classes: A possible model of social stereotyping. *Journal of the Experimental Analysis of Behavior*, *56*, 505–518.
- LESLIE, J., TIERNEY, K. J., ROBINSON, C. P., KEENAN, M., WATT, A., & BARNES, D. (1993). Differences between clinically anxious and nonanxious subjects in a stimulus equivalence training task involving threat words. *The Psychological Record*, *43*, 153–161.
- MCGLINCHY, A., KEENAN, M., & DILLENBURGER, K. (2000). Outline for the development of a screening procedure for children who have been sexually abused. *Research on Social Work Practice*, *10*, 721–747.
- MERWIN, R. M., & WILSON, K. G. (2005). Preliminary findings on the effects of self-referring and evaluative stimuli on stimulus equivalence class formation. *The Psychological Record*, *55*, 561–575.
- MITCHELL, C. J., DE HOUWER, J., & LOVIBOND, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences*, *32*, 183–198.
- MOXON, P. D., KEENAN, M., & HINE, L. (1993). Gender-role stereotyping and stimulus equivalence. *The Psychological Record*, *43*, 381–393.
- NEVIN, J. A. (1974). Response strength in multiple schedules. *Journal of the Experimental Analysis of Behavior*, *21*, 389–408.
- NEVIN, J. A., & GRACE, R. C. (2000). Behavioral momentum and the law of effect. *Behavioral and Brain Sciences*, *23*, 73–130.
- NOSEK, B. A., & BANAJI, M. R. (2001). The go/no-go association task. *Social Cognition*, *19*, 625–666.
- O'HORA, D., ROCHE, B., BARNES-HOLMES, D., & SMEETS, P. (2002). Response latencies to multiple derived stimulus relations: Testing two predictions of relational frame theory. *The Psychological Record*, *52*, 51–75.
- OLSON, M. A., & FAZIO, R. H. (2003). Relations between implicit measures of prejudice: What are we measuring? *Psychological Science*, *14*, 636–639.
- O'TOOLE, C., & BARNES-HOLMES, D. (2007). A derived transfer of functions and the Implicit Association Test. *Journal of the Experimental Analysis of Behavior*, *88*, 263–283.

- PAULHUS, D. L. (2002). Socially desirable responding: The evolution of a construct. In H. I. Braun, D. N. Jackson, & D. E. Wiley (Eds.), *The role of constructs in psychological and educational measurement* (pp. 46–69). Mahwah, NJ: Erlbaum.
- PILGRIM, C., & GALIZIO, M. (1996). Stimulus equivalence: A class of correlations or a correlation of classes? In T. R. Zentall & P. M. Smeets (Eds.), *Stimulus class formation in humans and animals* (pp. 173–195). Amsterdam, The Netherlands: Elsevier Science.
- RIDGEWAY, I., ROCHE, B., GAVIN, A., & RUIZ, M. R. (2010). Establishing and eliminating IAT effects in the laboratory: Extending a behavioral model of the Implicit Association Test. *European Journal of Behavior Analysis, 11*, 133–150.
- ROCHE, B., & BARNES, D. (1996). Arbitrarily applicable relational responding and sexual categorization: A critical test of the derived difference relation. *The Psychological Record, 46*, 451–475.
- ROCHE, B., BARNES, D., & SMEETS, P. (1997). Incongruous stimulus pairing and conditional discrimination training: Effects on relational responding. *Journal of the Experimental Analysis of Behavior, 68*, 143–160.
- ROCHE, B., LINHEHAN, C., WARD, T., DYMOND, S., & REHFELDT, R. (2004). The unfolding of the relational operant: A real-time analysis using electroencephalography and reaction time measures. *International Journal of Psychology and Psychological Therapy, 4*, 587–603.
- ROCHE, B., RUIZ, M., O'RIORDAN, M., & HAND, K. (2005). A relational frame approach to the psychological assessment of sex offenders. In M. Taylor & E. Quayle (Eds.), *Viewing child pornography on the Internet: Understanding the offence, managing the offender, and helping the victims* (pp. 109–125). Dorset, England: Russell House.
- ROTHERMUND, K., & WENTURA, D. (2004). Underlying processes in the Implicit Association Test (IAT): Dissociating salience from associations. *Journal of Experimental Psychology: General, 133*, 139–165.
- SIDMAN, M. (1960). *Tactics of scientific research*. New York, NY: Basic Books.
- SPENCER T. J., & CHASE, P. N. (1996). Speed analyses of stimulus equivalence. *Journal of the Experimental Analysis of Behavior, 65*, 643–659.
- STEELE, D., & HAYES, S. C. (1991). Stimulus equivalence and arbitrarily applicable relational responding. *Journal of the Experimental Analysis of Behavior, 56*, 519–555.
- STEFFENS, M. (2004). Is the Implicit Association Test immune to faking? *Experimental Psychology, 51*, 165–179.
- STEFFENS, M. C., KIRSCHBAUM, M., & GLADOS, P. (2008). Avoiding stimulus confounds in Implicit Association Tests by using the concepts as stimuli. *British Journal of Social Psychology, 47*, 217–243.
- STEFFENS, M., & PLEWE, I. (2001). Items' cross-category associations as a confounding factor in the implicit association test. *Zeitschrift Fuer Experimentelle Psychologie, 48*, 123–134.
- SMYTH, S., BARNES-HOLMES, D., & BARNES-HOLMES, Y. (2008). Acquired equivalence in human discrimination learning: The role of propositional knowledge. *Journal of Experimental Psychology: Animal Behavior Processes, 34*, 167–177.
- TYNDALL, I., ROCHE, B., & JAMES, J. E. (2004). The relationship between stimulus function and stimulus equivalence: A systematic investigation. *Journal of the Experimental Analysis of Behavior, 81*, 257–266.

- TYNDALL, I., ROCHE, B., & JAMES, J. E. (2009). The interfering effect of emotional stimulus functions on stimulus equivalence class formation: Implications for the understanding and treatment of anxiety. *European Journal of Behavior Analysis, 10*, 121–140.
- VERSCHUERE, B., PRATI, V., & DE HOUWER, J. (2009). Cheating the lie detector: Faking in the autobiographical Implicit Association Test. *Psychological Science, 20*, 410–413.
- WATT, A., KEENAN, M., BARNES, D., & CAIRNS, E. (1991). Social categorization and stimulus equivalence. *The Psychological Record, 41*, 33–50.
- WILLIAMS, J. M., MATHEWS, A., & MACLEOD, C. (1996). The emotional Stroop task and psychopathology. *Psychological Bulletin, 120*, 3–24.
- WILLIAMS, J. M., WATTS, F. N., MACLEOD, C., & MATHEWS, A. (1988). *Cognitive psychology and emotional disorders*. Chichester, England: Wiley.
- WULFERT, E., & HAYES, S. C. (1988). Transfer of a conditional ordering response through conditional equivalence classes. *Journal of the Experimental Analysis of Behavior, 50*, 125–144.