Spanish Journal of Psychology (2014), **17**, e73, 1–7.

doi:10.1017/sjp.2014.67

GRace: A MATLAB-Based Application for Fitting the Discrimination-Association Model

Luca Stefanutti, Michelangelo Vianello, Pasquale Anselmi and Egidio Robusto

University of Padua (Italy)

Abstract. The Implicit Association Test (IAT) is a computerized two-choice discrimination task in which stimuli have to be categorized as belonging to target categories or attribute categories by pressing, as quickly and accurately as possible, one of two response keys. The discrimination association model has been recently proposed for the analysis of reaction time and accuracy of an individual respondent to the IAT. The model disentangles the influences of three qualitatively different components on the responses to the IAT: stimuli discrimination, automatic association, and termination criterion. The article presents *General Race* (GRace), a MATLAB-based application for fitting the discrimination association model to IAT data. GRace has been developed for Windows as a standalone application. It is user-friendly and does not require any programming experience. The use of GRace is illustrated on the data of a Coca Cola-Pepsi Cola IAT, and the results of the analysis are interpreted and discussed.

Received 2 July 2013; Revised 15 January 2014; Accepted 24 January 2014

Keywords: implicit association test, discrimination association model, response time, accuracy, implicit measures.

The Implicit Association Test (IAT, Greenwald, McGhee, & Schwartz, 1998) is the most popular procedure for measuring automatic associations. Since its appearance, the IAT has received significant attention as an effective tool for the investigation of implicit social cognitions, including attitudes (Greenwald, Smith, Sriram, Bar-Anan, & Nosek, 2009; Maison, Greenwald, & Bruin, 2004; Uhlmann, Dasgupta, Elgueta, Greenwald, & Swanson, 2002), stereotypes (Cvencek, Meltzoff, & Greenwald, 2011; Nosek, Banaji, & Greenwald, 2002; White & White, 2006) and self-concept (Schnabel, Asendorpf, & Greenwald, 2008; Steffens & Schulze König, 2006; Yamaguchi et al., 2007).

The discrimination-association model (DAM, Stefanutti, Robusto, Vianello, & Anselmi, 2013) is a stochastic model for the analysis of reaction time and accuracy of an individual respondent to the IAT. Specific features of the model make it a useful tool for the analysis of IAT data. Since the DAM accounts for both response time and accuracy, it uses the complete information provided by the IAT. The DAM disentangles the influences of three qualitatively different components on the responses to the IAT: stimuli discrimination, automatic association, and termination criterion. These components have interesting interpretation and allow for substantive insights about the response process. In particular, termination criterion informs about task difficulty and individual cautiousness, stimuli discrimination informs about the functioning of the stimuli, automatic association informs about the association between targets and attributes. Applications of the DAM to empirical data have shown that the model enables a fine-grained analysis of the IAT (Anselmi, Vianello, Stefanutti, & Robusto, 2013; Stefanutti et al., 2013).

This article presents *General Race* (GRace), a MATLAB-based application for fitting the DAM to IAT data. GRace is user-friendly, and does not require any programming experience.

The next sections (1) briefly describe the IAT and (2) the DAM, (3) illustrate the use of GRace on the data of a Coca Cola-Pepsi Cola IAT, and (4) interpret and discuss the results of the analysis. Details about the DAM and the fitting procedure can be found in Stefanutti et al. (2013).

Implicit Association Test

The IAT is a computerized two-choice discrimination task. A pair of target categories (e.g., *Coca Cola* and *Pepsi Cola*) and a pair of attribute categories (e.g., *Good* and *Bad*) are displayed at the top-left and top-right screen corners. Stimuli representing each of the categories appear, one at a time, in the center of the computer screen, and participants have to categorize them into one of the categories by pressing, as quickly and accurately as possible, one of two response keys.

The IAT consists of seven blocks. Block 1 involves the categorization of stimuli representing the target categories, whereas Block 2 involves the categorization of

Correspondence concerning this article should be addressed to Luca Stefanutti. Department FISPPA. University of Padua. Via Venezia, 8. 35131. Padua (Italy). Phone: +39–0498276687. Fax: +39–0498276600.
E-mail: luca.stefanutti@unipd.it

stimuli representing the attribute categories. Block 3 involves the categorization of stimuli representing the target categories and stimuli representing the attribute categories, with a certain target-attribute pairing. For example, the categories *Coca Cola* and *Good* share a response key, and the categories *Pepsi Cola* and *Bad* share the other. Block 4 repeats this task with an additional set of trials. Block 5 repeats the task of Block 2, with a reversed position of the attribute categories on the screen. Blocks 6 and 7 reverse the target-attribute pairing of Blocks 3 and 4, so that the categories *Coca Cola* and *Bad* share a response key, and the categories *Pepsi Cola* and *Good* share the other (Blocks 3 and 4 are counterbalanced across participants with Blocks 6 and 7). Blocks 1, 2 and 5 are called *practice* blocks, whereas Blocks 3, 4, 6 and 7 are referred to as critical blocks. The target-attribute pairing that leads to faster and more accurate responses is called *compatible*, whereas the other is called *incompatible*.

Discrimination-Association Model

The DAM assumes that, when a stimulus is presented on the computer screen, the response of the participant is governed by four parallel and independent processes, one for each category of the IAT. In the Coca Cola-Pepsi Cola IAT at hand, these processes are denoted as $X_C(t)$, $X_P(t)$, $X_G(t)$, and $X_B(t)$, where C = *Coca Cola*, P = *Pepsi Cola*, G = *Good*, and B = *Bad*. Model assumptions are that every stimulus potentially contains - albeit in a variable quantity - evidence for each of the four categories, and that it is simultaneously and independently processed by each of the four processes. Once a stimulus is presented on the screen, each process starts accumulating, on its own counter, selective information about a specific characteristic of it (for instance, process $X_C(t)$ accumulates information about the membership of the stimulus to category *Coca Cola*). The process which accrues the required amount of information (called *termination criterion*) in the shortest time produces the observable response. The four processes are assumed to behave as Poisson processes. This implies that, in each process, interarrival times (i.e., time intervals between consecutive units of information) are independent and exponentially distributed with rate λ (see, e.g., Townsend & Ashby, 1983).

Model parameters are the rates at which information accumulates on the counter of each process, and the termination criteria. There are 16 different rates (λ parameters), one for each pair that can be formed by taking one of the four processes and one of the four categories. For $i, j \in \{C, P, G, B\}$, the parameter $\lambda_{i \to j}$ is the average amount of information provided, in the time unit, by a stimulus of category i to process X_j(t). The 16 rates can be grouped into discrimination rates and association rates.

The discrimination rates regard the amount of information that target (respectively attribute) categories accumulate when target (resp. attribute) stimuli are presented. The rates $\lambda_{C \to C}$, $\lambda_{P \to P}$, $\lambda_{G \to G}$, $\lambda_{B \to B}$ are involved in the correct discrimination of the stimuli, whereas the rates $\lambda_{C \to P}$, $\lambda_{P \to C}$, $\lambda_{G \to B}$, $\lambda_{B \to G}$ are involved in the incorrect discrimination. The discrimination rates provide information about stimuli discrimination, that is, whether the stimuli that have been chosen to represent a certain category are easily recognized and correctly categorized in their own category, rather than incorrectly categorized in the contrasted category. The better the discrimination, the smaller the incorrect discrimination rates compared with the correct discrimination rates. In a Flowers-Insects IAT, for example, the ratio between correct and incorrect discrimination rates of category *Flowers* is expected to be larger if the words used for representing the category refer to familiar flowers (e.g., *rose* and *tulip*) than if they refer to unfamiliar flowers (e.g., *hydrangea* and *zephyranth*).

The association rates regard the amount of information that target (resp. attribute) categories accumulate when attribute (resp. target) stimuli are presented. In particular, $\lambda_{C \to G}$, $\lambda_{C \to B}$, $\lambda_{P \to G}$, $\lambda_{P \to B}$ are the rates at which information concerning membership to attribute categories is accumulated when a target stimulus is presented (*target-driven* associations), and $\lambda_{G \to C}$, $\lambda_{G \to P}$, $\lambda_{B \to C}$, $\lambda_{B \to P}$ are the rates at which information concerning membership to target categories is accumulated when an attribute stimulus is presented (*attribute-driven* associations). The association rates express the amount of information that targets (resp. attributes) provide about attributes (resp. targets) and, in this sense, they can be seen as an expression of the association strenght between targets and attributes. The stronger the association, the greater the value of these parameters.

In practical applications of the model, the association rates might allow the identification of patterns of *automatic association* between targets and attributes that differ from one individual to another in both structure and meaning. For example, in the Coca Cola-Pepsi Cola IAT at hand, the association rates might allow the distinction between individuals whose implicit preference for a certain cola results from a positive evaluation of that cola from those whose implicit preference results from a negative evaluation of the constrasting cola. In a White-Black IAT, they might allow the distinction between individuals with an implicit outgroup derogation (e.g., Black-Bad association in a white participant) and individuals with an implicit ingroup favouritism (e.g., White-Good association in a white participant).

The termination criteria concern the amount of information that has to be accumulated before a response is given. Individuals are more cautious to the extent they perceive a block of trials as difficult

https://doi.org/10.1017/sjp.2014.67 Published online by Cambridge University Press

(Brendl, Markman, & Messner, 2001). The practice blocks involve the categorization of stimuli representing either the target categories or the attribute categories, whereas the critical blocks involve the categorization of stimuli representing the target categories and stimuli representing the attribute categories. As a consequence, the critical blocks are expected to be more difficult than the practice blocks. The critical blocks might differ in difficulty. For instance, an individual with a close association between *Flowers* and *Good* will find it easier to respond when *Flowers* shares the response key with *Good* than when it shares the response key with *Bad*. An incorrect response is expected to be the effect of carelessness or inattention. Therefore, the termination criteria of the wrong responses are expected to be lower than those of the correct responses.

The termination criteria may be the combined result of individual cautiousness, task difficulty, and their interaction. There are 6 termination criteria (*K* parameters), one for each pair that can be formed by considering a particular type of block (practice, compatible, incompatible) and a response category (correct, wrong).

GRace

GRace (availabe upon request from the corresponding author) is a free standalone application developed in MATLAB for Windows. The users of GRace will be provided with detailed information and instructions about the application.

It is possible to open the data file that has to be analyzed by selecting "Open data" from the "File" menu. MATLAB files (.mat), EXCEL files (.xls, .xlsx), and tab-delimited text files (.txt) are accepted. The data file should provide trials in rows and the following variables in columns: (1) participant code, (2) block code, (3) trial code, (4) response accuracy (0 = wrong, 1 = correct), and (5) response time in milliseconds (see Figure 1 for

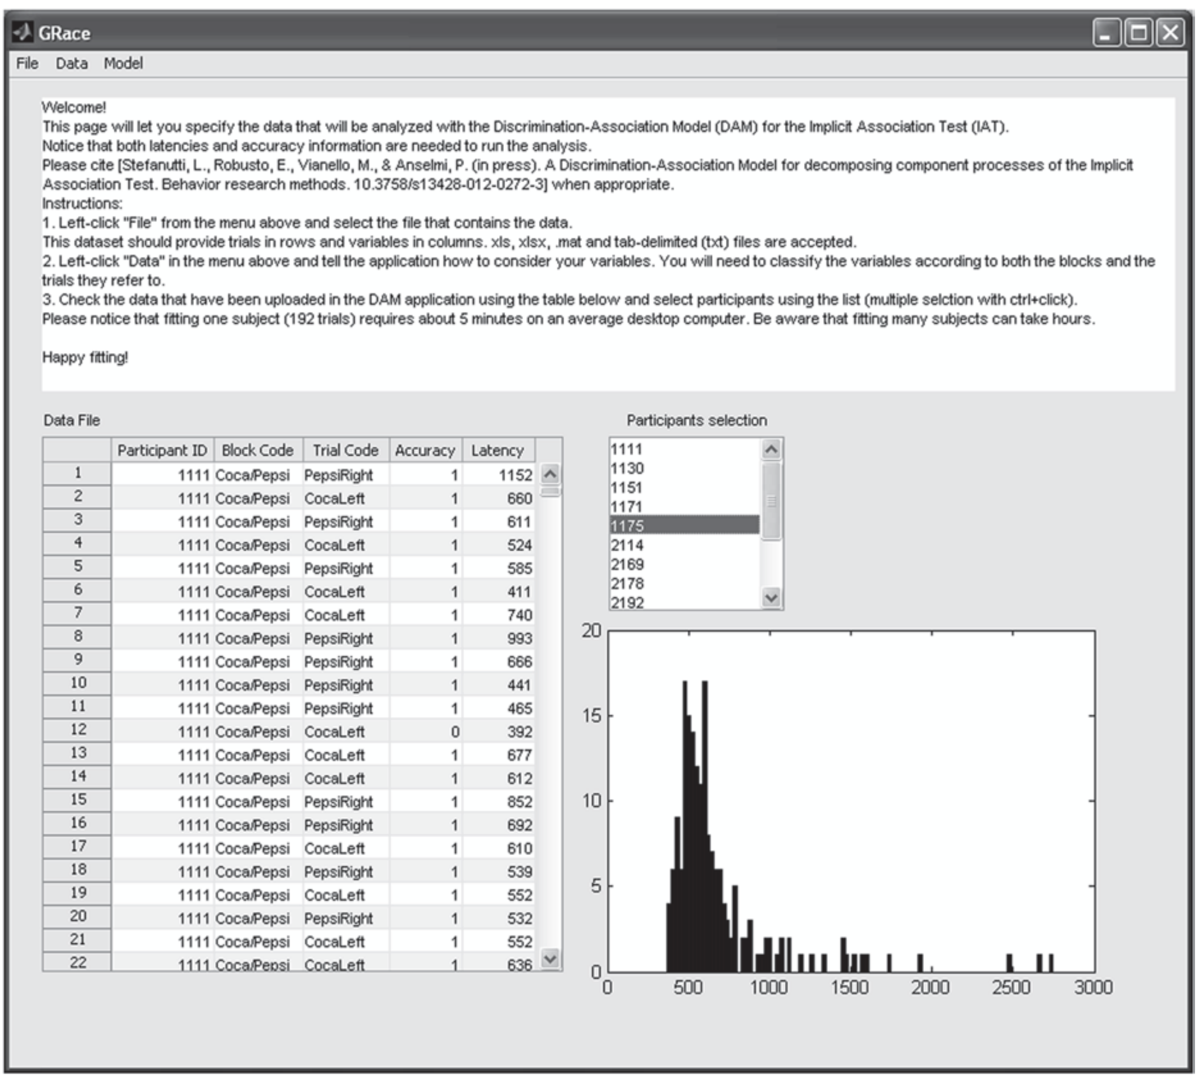

Figure 1. Screenshoot of the data and selection of the participants who have to be analyzed.

https://doi.org/10.1017/sjp.2014.67 Published online by Cambridge University Press

an example). In the IAT, the built-in error correction is typically used to inflate the latencies of the wrong responses with the time elapsed to correct them. The DAM applies to latencies without correction.

The "Participants selection" listbox (Figure 1) displays the codes of all participants that are present in the data file. It is possible to fit the DAM on subsets of participants by selecting their codes from this listbox. In our example, participant 1175 is selected. The histogram in Figure 1 depicts the response time distribution of the selected participant.

After this preliminary operation, it is necessary to specify and classify blocks and trials in the data file. Information about the blocks is provided by selecting "Classify block codes" from the "Data" menu. The "Available Block Codes" listbox (Figure 2) displays the codes that identify the different IAT blocks in the data file. Each block code is selected from the listbox and classified as practice, compatible or incompatible block. In our example, "Coca/Pepsi", "Bad/Good" and "Pepsi/Cola" are classified as practice blocks, "PepsiBad/CocaGood" and "PepsiBad/CocaGood_2" as compatible blocks, and "CocaBad/PepsiGood" and "CocaBad/PepsiGood_2" as incompatible blocks. Please note that, at this stage (i.e., before the data are analyzed), the distinction between compatible and incompatible blocks only reflects a conventional classification and will only affect the interpretation of the output.

Information about the trials is provided by selecting "Classify trial codes" from the "Data" menu. The listbox on the left displays the codes that identify the trials in the data file (Figure 3). Each trial code is selected from the listbox and classified into one of the four stimulus categories. In our example, "CocaLeft" and "CocaRight" are classified as Target 1, "PepsiLeft" and "PepsiRight" as Target 2, "GoodRight" as Attribute 1 and "BadLeft" as Attribute 2.

At this point, it is possible to fit the DAM by selecting "Fit" from the "Model" menu. Maximum likelihood estimation of model parameters is accomplished by the BFGS optimization algorithm (Broyden, 1970; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970), an iterative procedure that terminates when one of the following two conditions is reached and satisfited: (1) the maximum number of iterations is exceeded, or (2) the decrease of log-likelihood in two successive iterations is less than a sufficiently small tolerance value. By default, the two criteria are respectively set to 10,000 and 10^{-6}. By default, model parameters are estimated only once on the same data, and trials whose latencies are outliers in the response time distribution are

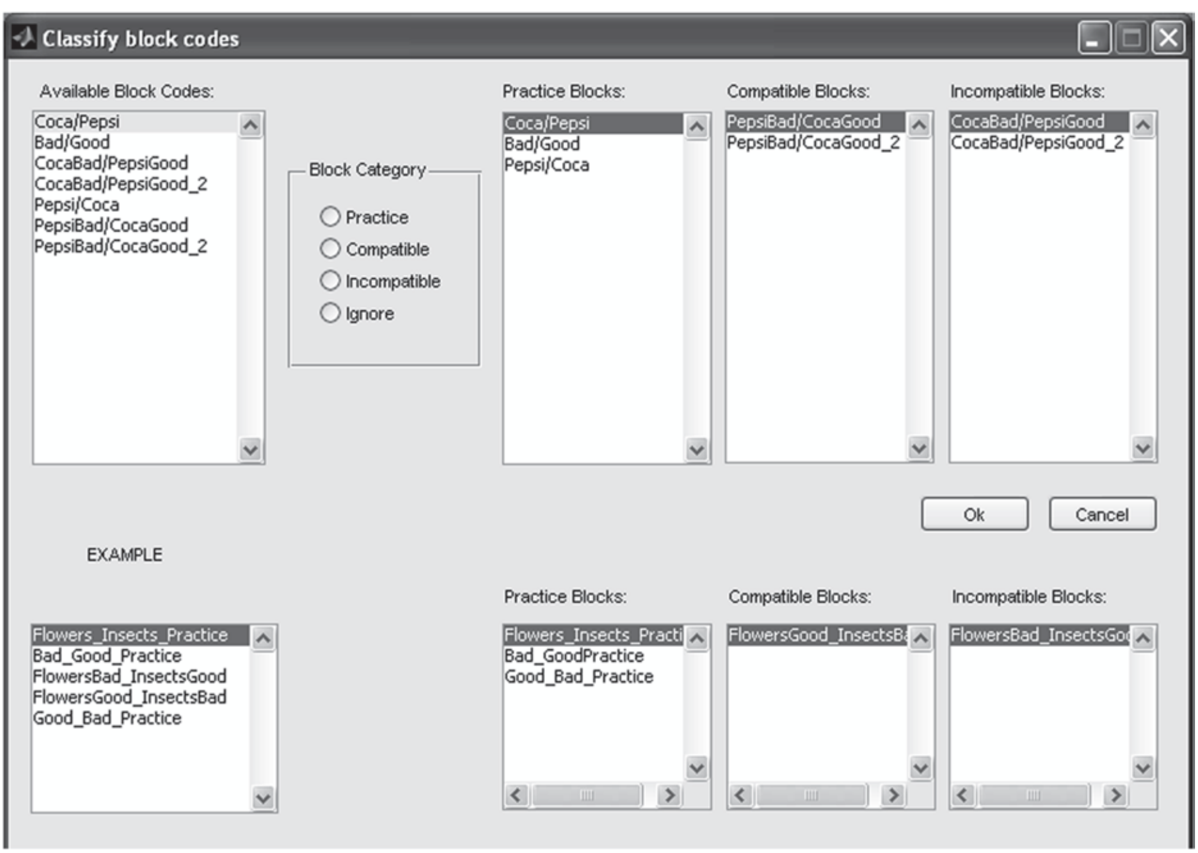

Figure 2. Classification of the block codes as practice, compatible and incompatible.

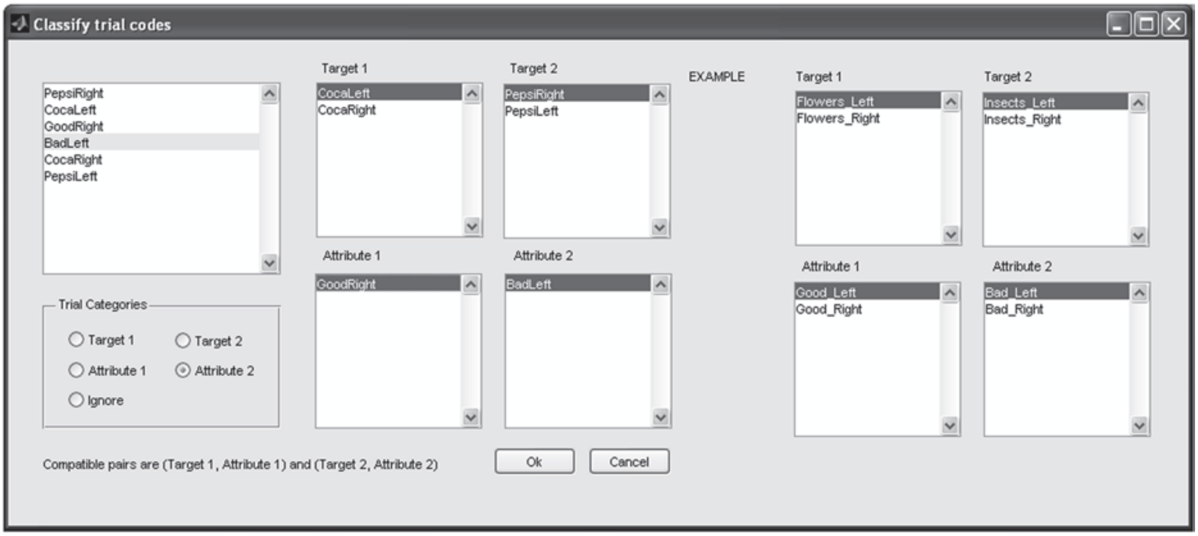

Figure 3. Classification of the trial codes as Target 1, Target 2, Attribute 1 and Attribute 2.

discarded according to Tukey's criterion (see, e.g., Hoaglin, Mosteller, & Tukey, 1983). All these default settings can be modified by selecting "Options" from the "Model" menu. In particular, estimating model parameters more than once is useful in an attempt to avoid local maxima. Only the solution with the largest likelihood is retained and displayed in the output file.

The output of GRace

The first part of the output (Figure 4) reports the percentage of trials discarded from the analysis (if outlier

```
Cutoff value for outlier latencies
      Minimum latency:    181 ms    (0.0% of 184)
      Maximum latency:   1014 ms    (10.3% of 184)

Discrimination-Association Model
      Exitflag:                   1
      Gradient norm:              0.00252616
      Condition number:           4815.94
      Log-likelihood:             5.42785
      AIC:                        54.8557
      Chi-square:                 33.6132
      Degrees of freedom:         39
      P-value:                    0.713638
      Minimum expected frequency  2

LAMBDA PARAMETERS ESTIMATION

lambda                                              estimate        SE      lb95%      ub95%
Target 1 -> Target 1 correct discrimination           24.063     0.032     23.453     24.681
Target 2  -> Target 1 incorrect discrimination         1.227     0.339      0.197      3.139
Target 1  -> Target 2 incorrect discrimination         0.804     0.573      0.004      4.075
Target 2  -> Target 2 correct discrimination          24.292     0.032     23.674     24.918
Attrib. 1 -> Attrib. 1 correct discrimination         21.487     0.034     20.881     22.103
Attrib. 2 -> Attrib. 1 incorrect discrimination        0.450     0.878      0.001      5.740
Attrib. 1 -> Attrib. 2 incorrect discrimination        1.666     0.306      0.478      3.571
Attrib. 2 -> Attrib. 2 correct discrimination         19.796     0.035     19.196     20.405
Target 1 -> Attrib. 1 association                      5.817     0.165      4.360      7.483
Target 2 -> Attrib. 1 association                      1.805     0.246      0.743      3.332
Target 1 -> Attrib. 2 association                      0.000   507.633    253.071 1294613.046
Target 2 -> Attrib. 2 association                      4.217     0.232      2.556      6.291
Attrib. 1 -> Target 1 association                      5.861     0.169      4.364      7.579
Attrib. 2 -> Target 1 association                      3.834     0.129      2.909      4.886
Attrib. 1 -> Target 2 association                      0.000   203.910     40.834 208890.153
Attrib. 2 -> Target 2 association                      5.990     0.153      4.616      7.543


KAPPA PARAMETERS ESTIMATES

kappa                        estimate        SE      lb95%      ub95%
Practice (corr. resp.)         25.010     0.033     24.385     25.644
Practice (wrong resp.)          4.468     0.261      2.826      6.632
Compatible (corr. resp.)       31.678     0.036     30.896     32.471
Compatible (wrong resp.)        9.137     0.127      7.783     10.614
Incompatible (corr. resp.)     34.130     0.033     33.387     34.882
Incompatible (wrong resp.)     14.135     0.123     12.444     15.942
```

Figure 4. The output of Grace.

https://doi.org/10.1017/sjp.2014.67 Published online by Cambridge University Press

deletion has been applied) and a number of indices and statistics concerning the fitting procedure (information about the meaning of the indices is also provided). As a goodness-of-fit index of the DAM, a Pearson's Chi-square statistic is computed, that is based on the procedure described by Klauer, Voss, Schmitz, & Teige-Mocigemba (2007). In our example, we can see that the model adequately fits the data of the participant at hand ($\chi^2(39) = 33.61$, $p = .71$).

The following part of the output reports the estimates of the λ parameters. The labels on the left of the arrow denote the stimulus category, whereas the labels on the right denote the process. It is worth recalling that, in the present example, Target 1, Target 2, Attribute 1 and Attribute 2 respectively denote *Coca Cola*, *Pepsi Cola*, *Good* and *Bad*.

The first eight λ parameters are discrimination rates. For each stimulus category, the rate concerning the correct discrimination was greater than that concerning the incorrect discrimination ($\lambda_{\text{Target1}\rightarrow\text{Target1}} > \lambda_{\text{Target1}\rightarrow\text{Target2}}$; $\lambda_{\text{Target2}\rightarrow\text{Target2}} > \lambda_{\text{Target2}\rightarrow\text{Target1}}$; $\lambda_{\text{Attrib.1}\rightarrow\text{Attrib.1}} > \lambda_{\text{Attrib.1}\rightarrow\text{Attrib.2}}$; $\lambda_{\text{Attrib.2}\rightarrow\text{Attrib.2}} > \lambda_{\text{Attrib.2}\rightarrow\text{Attrib.1}}$; the parameters can be standardized and their statistical difference can be tested). The stimuli provided, in the time unit, more information towards the correct response than towards the incorrect response. This means that the participant at hand easily recognized the stimuli and categorized them into the correct category.

The last eight λ parameters are association rates. The target-driven association rates are considered first. The *Coca Cola* stimuli provided, in the time unit, more information towards the category *Good* than towards the category *Bad* ($\lambda_{\text{Target1}\rightarrow\text{Attrib.1}} > \lambda_{\text{Target1}\rightarrow\text{Attrib.2}}$), whereas the *Pepsi Cola* stimuli provided more information towards the category *Bad* than towards the category *Good* ($\lambda_{\text{Target2}\rightarrow\text{Attrib.2}} > \lambda_{\text{Target2}\rightarrow\text{Attrib.1}}$). Accumulation of information about the category *Good* was faster when the stimuli were *Coca Cola* and slower when the stimuli were *Pepsi Cola* ($\lambda_{\text{Target1}\rightarrow\text{Attrib.1}} > \lambda_{\text{Target2}\rightarrow\text{Attrib.1}}$). Accumulation of information about the category *Bad* was faster when the stimuli were *Pepsi Cola* and slower when the stimuli were *Coca Cola* ($\lambda_{\text{Target2}\rightarrow\text{Attrib.2}} > \lambda_{\text{Target1}\rightarrow\text{Attrib.2}}$). The same pattern of results is obtained by considering the attribute-driven association rates. In this participant, an implicit preference for *Coca Cola* relative to *Pepsi Cola* is observed, that resulted from both a positive evaluation of the *Coca Cola* and a negative evaluation of the *Pepsi Cola*. Different patterns of automatic association have been described, in which the preference for a certain cola either resulted only from a positive evaluation of that cola or resulted only from a negative evaluation of the contrasted cola (Stefanutti et al., 2013).

The standard errors of two λ parameters are very large. A possible explanation is that the data at hand are not sufficiently informative for computing reliable estimates of the two parameters. Empirical nonidentification of one or more parameters might occur even if the model is identifiable in theory (see Stefanutti et al., 2013). Another possible explanation is that the true values of the two parameters are on or too close to the boundary value of 0.

The last part of the output reports the estimates of the K parameters. It is worth recalling that, in the present example, the label *compatible* denotes the *Pepsi Cola-Bad/Coca Cola-Good* blocks and the label *incompatible* denotes the *Coca Cola-Bad/Pepsi Cola-Good* blocks. For the participant at hand, the *Pepsi Cola-Bad/Coca Cola-Good* blocks were more difficult than the practice blocks and less difficult than the *Coca Cola-Bad/Pepsi Cola-Good* blocks ($K_{\text{Practice}} < K_{\text{Compatible}} < K_{\text{Incompatible}}$). This holds for both the correct responses and the incorrect responses.

Conclusions

The article presented GRace, a MATLAB-based application for fitting the DAM to IAT data. The use of GRace has been illustrated on the data of a Coca Cola-Pepsi Cola IAT, and the results of the analysis have been interpreted and discussed.

Grace is constantly under development. Current work is devoted to allow users to impose equality constraints between the parameters, and to provide them with a formal test of the empirical identifiability of the parameters (for details, see Stefanutti et al., 2013). Theoretical extensions of the DAM are under development that incorporate nondecision components (e.g., encoding stimuli, motor response, distractions) within the model, and that allow for variability of rates and termination criteria within the blocks.

References

Anselmi P., Vianello M., Stefanutti L., & Robusto E. (2013). A poisson race model for the analysis of the Implicit Association Test. *TPM - Testing, Psychometrics, Methodology in Applied Psychology*, *20*, 249–261. http://dx.doi.org/10.4473/TPM20.3.4

Brendl C. M., Markman A. B., & Messner C. (2001). How do indirect measures of evaluation work? Evaluating the inference of prejudice in the Implicit Association Test. *Journal of Personality and Social Psychology*, *81*, 760–773. http://dx.doi.org/10.1037/0022-3514.81.5.760

Broyden C. G. (1970). The convergence of a class of double-rank minimization algorithms. *Journal of the Institute of Mathematics and its Applications*, *6*, 76–90. http://dx.doi.org/10.1093/imamat/6.1.76

Cvencek D., Meltzoff A. N., & Greenwald A. G. (2011). Math-gender stereotypes in elementary-school children. *Child Development*, *82*, 766–789. http://dx.doi.org/10.1111/j.1467-8624.2010.01529.x

Fletcher R. (1970). A new approach to variable metric algorithms. *Computer Journal*, *13*, 317–322. http://dx.doi.org/10.1093/comjnl/13.3.317

Goldfarb D. (1970). A family of variable metric updates derived by variational means. *Mathematics of Computing*, *24*, 23–26.

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. http://dx.doi.org/10.1037/0022-3514.74.6.1464

Greenwald A. G., Smith C. T., Sriram N., Bar-Anan Y., & Nosek B. A. (2009). Race attitude measures predicted vote in the 2008 U. S. Presidential Election. *Analyses of Social Issues and Public Policy*, *9*, 241–253. http://dx.doi.org/10.1111/j.1530-2415.2009.01195.x

Hoaglin D. C., Mosteller F., & Tukey J. W. (1983). *Understanding robust and exploratory data analysis*. (Eds.), New York, NY: John Wiley & Sons.

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. http://dx.doi.org/10.1037/0022-3514.93.3.353

Maison D., Greenwald A. G., & Bruin R. H. (2004). Predictive validity of the Implicit Association Test in studies of brands, consumer attitudes, and behavior. *Journal of Consumer Psychology*, *14*, 405–415. http://dx.doi.org/10.1207/s15327663jcp1404_9

Nosek B. A., Banaji M. R., & Greenwald A. G. (2002). Math = male, me = female, therefore math ≠ me. *Journal of Personality and Social Psychology*, *83*, 44–59. http://dx.doi.org/10.1037//0022-3514.83.1.44

Schnabel K., Asendorpf J. B., & Greenwald A. G. (2008). Implicit Association Tests: A landmark for the assessment of implicit personality self-concept. In G. J. Boyle, G. Matthews, & D. H. Saklofske (Eds.), *Handbook of Personality Theory and Testing* (pp. 508–528). London, UK: Sage.

Shanno D. F. (1970). Conditioning of quasi-Newton methods for function minimization. *Mathematics of Computing*, *24*, 647–656. http://dx.doi.org/10.2307/2004840

Stefanutti L., Robusto E., Vianello M., & Anselmi P. (2013). A discrimination-association model for decomposing component processes of the Implicit Association Test. *Behavior Research Methods*, *45*, 393–404. http://dx.doi.org/10.3758/s13428-012-0272-3

Steffens M. C., & Schulze König S. (2006). Predicting spontaneous big five behavior with Implicit Association Tests. *European Journal of Psychological Assessment*, *22*, 13–20. http://dx.doi.org/10.1027/1015-5759.22.1.13

Townsend J. T., & Ashby F. G. (1983). *Stochastic modeling of elementary psychological processes*. New York, NY: Cambridge University Press.

Uhlmann E., Dasgupta N., Elgueta A., Greenwald A. G., & Swanson J. E. (2002). Subgroup prejudice based on skin color among Hispanics in the United States and Latin America. *Social Cognition*, *20*, 198–226. http://dx.doi.org/10.1521/soco.20.3.198.21104

White M. J., & White G. B. (2006). Implicit and explicit occupational gender stereotypes. *Sex Roles*, *55*, 259–266. http://dx.doi.org/10.1007/s11199-006-9078-z

Yamaguchi S., Greenwald A. G., Banaji M. R., Murakami F., Chen D., Shiomura K., … Krendl A. (2007). Apparent universality of positive implicit self-esteem. *Psychological Science*, *18*, 498–500. http://dx.doi.org/10.1111/j.1467-9280.2007.01928.x