Journal of Contextual Behavioral Science 4 (2015) 157–162

Contents lists available at ScienceDirect

Journal of Contextual Behavioral Science

journal homepage: www.elsevier.com/locate/jcbs

Interpreting and inverting with less cursing: A guide to interpreting IRAP data

Ian Hussey [a,*,1], Miles Thompson [b], Ciara McEnteggart [a], Dermot Barnes-Holmes [c], Yvonne Barnes-Holmes [a]

[a] Maynooth University, Ireland
[b] Goldsmiths, University of London, United Kingdom
[c] Ghent University, Belgium

ARTICLE INFO

Article history:
Received 10 October 2014
Received in revised form
5 March 2015
Accepted 14 May 2015

Keywords:
IRAP
Implicit Relational Assessment Procedure

ABSTRACT

This Professional Interest Brief seeks to provide a clear guide to interpreting data generated by Implicit Relational Assessment Procedure (IRAP). The interpretation of IRAP data is not immediately intuitive and yet has received little explicit attention in the published literature. As such, it is hoped that this guide will help clarify this matter, particularly for those new to using the IRAP or intending to use the measure in the future. In doing so, we hope to make the measure more accessible and facilitate continued use of the methodology and its contribution to the contemporary Relational Frame Theory (RFT) literature.

1. Introduction

One of the cornerstones of Contextual Behavioral Science (CBS: Hayes, Barnes-Holmes, & Wilson, 2012) is its appeal to a basic account of human language and cognition through Relational Frame Theory (RFT). Relational Frame Theory argues that the fundamental building block of human cognitive abilities, such as abstract reasoning and generative language is "arbitrarily applicable relational responding" (AARR: see Hayes, Barnes-Holmes, and Roche, 2001). Much early RFT research revolved around demonstrating its proposed analytic units, relational frames, that were established in the laboratory (see Hughes and Barnes-Holmes, in press, for review). However, in recent years, RFT researchers have attempted to extend RFT's conceptual account by also assessing histories of relational responding that were established outside of the laboratory (see Barnes-Holmes, Barnes-Holmes, and Hussey, in press; Hussey, Barnes-Holmes, and Barnes-Holmes, 2015), such as by posing questions about the probability or "strength" of individuals' relational responding in applied domains such as obsessive compulsive tendancies, depression, or professional burnout (see Nicholson and Barnes-Holmes, 2012a, Hussey and Barnes-Holmes, 2012, Kelly and Barnes-Holmes, 2013, respectively). In order to do this, RFT researchers have built on methodologies frequently used in cognitive and social psychology to assess what are referred to as "implicit attitudes" (see De Houwer and Moors, 2010; see also Hughes, Barnes-Holmes, and Vahey, 2012). This has produced a procedure that has shown utility in assessing the relative strength of relational responding: the Implicit Relational Assessment Procedure (IRAP: Barnes-Holmes, Barnes-Holmes, Stewart, & Boles, 2010). IRAP research now represents one of the forefronts of RFT research (Barnes-Holmes et al., in press).

2. Task structure

A brief description of the procedure will now be provided, as the interpretation of IRAP data is best understood through an understanding of the structure of the task itself. The IRAP involves presenting pairs of stimuli to participants on a computer screen. Participants respond to blocks of these stimulus pairings, and are required to respond as accurately and quickly as possible according to what we will describe as two responding rules. In some IRAP studies, specific instructions regarding these rules are provided before each block (e.g., "Respond as if I am positive and others are negative"). However, in other studies, specific instructions to respond according to a particular rule for each block are not provided (e.g., "Try to get as many correct as possible – go fast, making a few errors is ok"). For the purposes of communication, however, we will describe the task in terms of utilizing two specific types of rule. In short, the IRAP compares, under accuracy and latency pressure, the relative ease (i.e., speed) with which participants

* Corresponding author.
E-mail address: Ian.Hussey@nuim.ie (I. Hussey).
[1] IH was assisted by a Government of Ireland Scholarship by the Irish Research Council.

http://dx.doi.org/10.1016/j.jcbs.2015.05.001
2212-1447/

Table 1
Example rules and stimuli for an IRAP investigating self-esteem.

Rule A	
Respond as if I am positive and others are negative	
Rule B	
Respond as if I am negative and others are positive	
Label 1: "Self"	**Label 2: "Others"**
I am	Other people are
I'm	Others are often
I think I am	Other people can be
Target 1: "Positive"	**Target 2: "Negative"**
Loyal	Manipulative
Trustworthy	Dishonest
Kind	Cruel
Moral	Horrible
Generous	Selfish
Friendly	Heartless
Response Option 1	**Response Option 2**
True	False

respond according to one rule relative to the other. In other words, the IRAP is a procedure that is used to assess subtle reaction time biases that are often referred to as reflecting "implicit attitudes" (De Houwer & Moors, 2010).

For illustrative purposes, consider the stimulus set for the hypothetical self-esteem IRAP outlined in Table 1 (see Vahey, Barnes-Holmes, Barnes-Holmes, and Stewart, 2010, for an alternative published version of a self-esteem IRAP). One of the two responding rules, Rule A (e.g., "Respond as if I am positive and others are negative") or Rule B (e.g., "Respond as if I am negative and others are positive"), is presented to participants before each block of trials.We will refer to these as Rule A blocks and Rule B blocks. Table 1 also lists what are arbitrarily referred to as label stimuli (presented at the top of the screen), target stimuli (presented in the middle of the screen) and response options (presented at the bottom of the screen). Label stimuli frequently contain what can loosely be referred to as categories (e.g., self vs. others), whereas target stimuli frequently contain attributes (e.g., positive vs. negative). These four classes of stimuli each contain one or more exemplars of the relevant category and attribute (e.g., loyal, trustworthy, kind, etc.). However, when describing the data researchers typically refer only to the overarching functional class (e.g., positive or self). Finally, participants respond using one of the two response options (e.g., same and different, or true and false), which are typically mapped to the "D" and "K" keys.

Each IRAP trial presents one label stimulus and one target stimulus and both response options. The combination of two label categories (e.g., self and others) and two target categories (e.g., positive and negative) produce four possible "trial-types" (e.g., trial-type 1=self-positive, trial-type 2=self-negative, trial-type 3=others-positive, and trial-type 4=others-negative). It is important to note that the trial-types are procedurally separated, insofar as label 1 stimuli are never presented within the same trial as label 2 stimuli, and target 1 stimuli are never presented within the same trial as target 2 stimuli.

The required correct and incorrect response options for each trial-type in each of the two rule blocks are pre-determined by the task structure itself (Table 2; see Barnes-Holmes, Barnes-Holmes, Stewart and Boles, 2010). To illustrate, let us return to the hypothetical self-esteem IRAP. Rule A block employs contingencies that require participants to respond as if "I am positive and others are negative". For example, a self-positive trial (i.e., trial-type 1) might present the participant with the stimuli "I am" and "loyal" and the response options "True" and "False". In this case, True would be the correct response, by definition, while selecting False would present the participant with a red X. However, if these same stimuli appeared on a Rule B block trial, the correct response would now be False. The IRAP is arranged in this way in order to assess the difference in reaction times between Rule A and Rule B blocks for each trial-type (e.g., the difference in speed between responding True on Rule A blocks vs. False on Rule B blocks). Furthermore, participants are presented with pairs of Rule A and Rule B blocks, each of which contains a large number of IRAP trials in order to capture a sufficient number of reaction times to conduct a meaningful analysis (e.g., 48). Typically, participants complete pairs of practice blocks until they meet both accuracy and latency mastery criteria, followed by three test block pairs (see Barnes-Holmes, Barnes-Holmes, Stewart, and Boles, 2010). Given that the IRAP effect is produced via accuracy and latency pressure, these criteria should be set as high as is feasible. Recent studies have frequently employed accuracy $\geq$ 80% and median time to first correct response $\leq$ 2000 ms, but future work may of course tighten these criteria further. It should be noted that both mastery criteria must be met within both blocks in a block pair for the criteria to have been met. On balance, variations on these criteria have not been systematically explored, and future efforts might revise these practices.

3. Interpretation of IRAP effects

3.1. Methods of quantifying effects on the IRAP

To reiterate, the IRAP presents stimuli to participants in pairs of blocks. The same categories of stimuli are presented in both blocks. However, the critical difference between the two blocks is that the required response option for each trial-type alternates between them. For example, on one block, participants must respond to a given stimulus pair (e.g., "I am" and "Loyal") with one response option (e.g., "True"), whereas on the other block, participants must respond with the other response option (e.g., "False"). The IRAP researcher then seeks to quantify the difference in responding speed between the two blocks in any pair. Loosely, this difference indicates which responding direction makes more intuitive sense or is more "automatic" for an individual (De Houwer & Moors, 2012).

While this difference can be quantified in numerous ways, specific common practices have emerged from the broader literature on the analysis of reaction-time data (see Balota and Yap, 2011; Ratcliff, 1993, Whelan, 2008). In particular, due to the distribution of reaction times, some form of normalization technique is recommended when quantifying the differences between block pairs. The most common way to quantify the difference between

Table 2
Required responses on each trial-type within a hypothetical self-esteem IRAP.

	Trial-type 1: self-positive	**Trial-type 2: self-negative**	**Trial-type 3: others-positive**	**Trial-type 4: others-negative**
Rule A block ("Respond as if I am positive and others are negative")	True	False	False	True
Rule B block ("Respond as if I am negative and others are positive")	False	True	True	False