

Modelling the IAT: Implicit Association Test Reflects Shallow Linguistic Environment and not Deep Personal Attitudes

Dermot Lynott (dermot.lynott@manchester.ac.uk),

Himanshu Kansal (himanshu.kansal@postgrad.mbs.ac.uk)

Decision and Cognitive Sciences Research Centre, Manchester Business School, University of Manchester
Booth Street West, Manchester M15 6PB, UK

Louise Connell (louise.connell@manchester.ac.uk)

School of Psychological Sciences, University of Manchester, Oxford Road, Manchester M13 9PL, UK

Kerry O'Brien (kerry.o'brien@monash.edu)

Behavioural Studies, Monash University, PO Box 197, Caulfield East VIC 3145, AUSTRALIA
and School of Psychological Sciences, University of Manchester, Oxford Road, Manchester M13 9PL, UK.

Abstract

People often have thoughts, attitudes and biases that are not themselves consciously aware of or that they would rather not share with others. To assess such attitudes, researchers use paradigms like the Implicit Association Test (IAT) that do not rely on explicit responding to determine the level of bias a person holds towards a particular target concept (e.g., race, gender, age). Responses in the IAT are assumed to reflect deeply held beliefs and attitudes, and not shallow, superficial associations. However, as linguistic distributional information has been shown to serve as a viable heuristic in many cognitive tasks, we investigated whether it could be used to predict the level of bias established by the IAT. We used a large corpus of language (Web 1T) and data from 16 IAT studies (N = 1825) to examine whether the degree of linguistic co-occurrence for target concepts and attributes reflected the size of bias observed in human behavioural data. We found that the effect size of the linguistic biases corresponded strongly with the effect sizes from the behavioural data. We suggest that language reflects prevalent cultural attitudes which are captured by tasks such as the IAT, suggesting that the IAT may reflect shallow, linguistic associations rather than deeper conceptual processing.

Keywords: linguistic distributional information; implicit association test; IAT; attitudes; model.

Introduction

If we openly asked people questions like "are you sexist" or "are you racist", we would probably expect people to be reluctant to respond, if we got any response at all. When asking for judgements on controversial topics and divisive issues, people have a strong desire to provide socially acceptable responses that may be contrary to their true beliefs (e.g., Furnham, 1986; Paulus, 1991). As such, there is often a disconnect between what people say and what they do. In order to avoid tasks that require explicitly thinking about a particular issue or that permit strategic responding by participants, researchers in social cognition have instead developed paradigms that try to tap into people's attitudes in a more implicit manner (Fazio, Jackson, Dunton & Williams, 1995; Greenwald, McGee & Schwartz, 1998). The most frequently used of these paradigms is the

Implicit Association Test, or IAT. The IAT is essentially a categorisation task, similar to many priming paradigms used across the cognitive sciences, designed to capture the degree of bias or prejudice that an individual has towards a particular concept. (e.g, race, age). We describe the task in more detail below.

A search using Google Scholar reveals that the IAT is referenced in over 4000 papers in the last 10 years alone. In spite of its widespread use, there is ongoing disagreement regarding what the IAT is actually measuring (Blanton et al., 2009; Fazio & Olsen, 2003; Greenwald, Poehlman, Uhlman & Banaji, 2009). Nonetheless, the creators and most proponents of the paradigm maintain that "the IAT assesses the strengths of associations between concepts" (p18., Greenwald, Poehlman, Uhlman & Banaji, 2009) and it is assumed to reflect deep, underlying, unconscious biases. However, when one situates the IAT within the broader context of cognitive research examining the structure of the conceptual system, such a claim is ambiguous.

Several researchers have described the conceptual system as comprising two distinct but interrelated components; a linguistic system and a simulation system (Barsalou, Santos, Simmons & Wilson, 2008; Connell & Lynott, 2011; 2012; Louwerse & Jeunieux, 2008). The linguistic system reflects language usage, and captures the distributional patterns (or statistical regularities) of words and phrases, making this system best suited for "quick and dirty" heuristic processing (Lynott & Connell, 2010). The simulation system, on the other hand, captures perceptual, affective and motor information from our environmental experience and is better suited to deep, slow, precise processing. Thus, performance in the IAT may reflect responses from one of these two systems, raising two alternative hypotheses. The first, is that the IAT indeed reflects personal attitudes emerging from deep-rooted, affective and conceptual processing in the simulation system. For example, an intelligence/obesity prejudice (O'Brien et al, 2007) would take the form of conceptual retrieval of a "fat person" automatically evoking associated concepts of "stupidity" and a negatively valenced affective association of "badness". This perspective is summarised by Nosek, Banaji and Greenwald (p112., 2002)

who argue that implicit attitudes "reveal the *deep* influence of the immediate environment and the broader culture on internalized preferences and beliefs".

The second option is that IAT scores reflect much shallower processing of the socio-cultural environment, specifically the token-to-token statistical patterns of the linguistic system. For example, people may often encounter the word "fat" in close proximity to the word "stupid" in conversations they hear or in texts they read, resulting in the automatic activation of the word "stupid" every time the word "fat" is encountered. Importantly, activating a word like "stupid" does not require full conceptual retrieval (Louwerse & Connell, 2011). Rather, linguistic associations like these operate at a shallow, superficial level that can produce a response to a given task without recourse to deeper conceptual or affective processing.

There is good reason to believe that the IAT may reflect the latter shallow linguistic associations rather than the former deeper, affective, conceptual attitudes. Several studies have shown that the linguistic system is used as a shortcut to provide a "good enough" response to conceptual tasks, whenever possible (e.g., Connell & Lynott, 2012; Louwerse & Connell, 2011). In particular, when processing demands are shallow and the participant is placed under time pressure the linguistic system provides a useful heuristic for responding without recourse to the greater computational expense of full, perceptual, affective and motor simulation of the concept. For example, conceptual tasks such as property-verification (e.g., making true/false judgments regarding object properties - *apple can be green*) can be successfully completed solely on the basis of the word-to-word associations of "apple" and "green"; these words frequently appear in close proximity and therefore it is a reasonable heuristic to assume that this property belongs to this concept. When participants respond quickly (Louwerse & Connell, 2011) or when the set of items is poorly constructed (Solomon & Barsalou, 2004) their responses are based on these linguistic associations and not on deeper conceptual representations. For example, using response time data from a property-verification task, Louwerse and Connell (2011) demonstrated that measures of distributional patterns from the linguistic system could be used to predict the faster responses of participants, but not their slower responses. Conversely, measures of the simulation system could predict slower responses, but not faster responses. Evidence that the IAT does not engage deeper processing is provided by a recent study by Foroni and Semin (2012). Foroni and Semin had two groups complete the IAT; one group completed the task as normal, while the other completed the task with facial feedback being inhibited by holding a pen between the lips during the task. Holding the pen in this position leads to sustained activation of the zygomaticus major muscle (used in frowning), which is normally activated following the presentation of a valenced stimulus. However, inhibition of this muscle during the IAT made no difference to the level of bias observed. This suggests that IAT does not engage

the affective system in a way that would be expected if the task required processing in the simulation system.

Given that the linguistic system is capable of providing quick and dirty responses in a variety of seemingly complex tasks and given that responses in the IAT may be of a superficial nature (i.e., not requiring the deeper processing of the simulation system), we considered whether IAT biases could be predicted by the statistical distributional associations in language. While the linguistic associations and simulation systems are closely related, they are not exact replications of each other because each system gains experience from a different source. Just because two words share a linguistic association in the socio-cultural environment, because they are sometimes juxtaposed, it does not mean that their referent concepts are tightly bound in a personal, affective/conceptual attitude. If IAT responses are predicted by linguistic associations then it suggests that the IAT itself is a shallow measure of the language structure to which an individual has been exposed and not necessarily a reflection of deeper biases. We describe below the IAT paradigm in more detail before outlining the current study.

Condition	Word belongs to...?	"stupid"	"smart"
Congruent	Bad OR Fat	fast	[incorrect]
	Good OR Thin	[incorrect]	fast
Incongruent	Bad OR Thin	slow	[incorrect]
	Good OR Fat	[incorrect]	slow

Table 1: Schematic of response patterns in an IAT on obesity prejudice. The first column describes category as congruent or incongruent pairings. The second column indicates the two judgements participants must consider for each target word. Third and fourth columns describe the predicted patterns for two target concepts "stupid" and "smart". "Incorrect" indicates a wrong answer (e.g., "stupid" should not be "good" or "thin")

The Implicit Association Test

The IAT represents one of the most frequently used paradigms for examining implicit attitudes (e.g., Greenwald, McGee & Schwartz, 1998), with hundreds of studies already published using this approach (see e.g., Greenwald, Poehlman, Uhlman & Banaji, 2009). The IAT is used to give an insight into people's automatically activated biases and prejudices and is designed to overcome the issues of strategising and socially desirable responding by participants. The IAT achieves this by requiring extremely rapid and accurate responses from participants to tap into automatic associations between some target concept and an attribute. For example, O'Brien and colleagues (O'Brien, Hunter & Banks, 2007) examined people's anti-fat prejudices using the IAT to see whether people associated obesity with negative concepts like stupidity. The IAT contrasts performance for a *congruent* pairing of targets and attributes (e.g., obesity-bad; thinness-good) to an

incongruent pairing of targets and attributes (e.g., obesity-good; thinness-bad). The participant's task is to categorise target stimuli as they appear on screen using one of these two pairings. In a congruent block, if the word "stupid" appeared onscreen, the participant would press the key indicating they belonged to the "fat OR bad" category, while if the word "smart" appeared they would press the key to categorise it as belonging to the "thin OR good" category. In this way, each target attribute has an identifiably correct response. Table 1 presents a schematic of the responses in both congruent and incongruent conditions. Every participant completes both categorisation pairings in a counterbalanced fashion.

The key question is, which pairing do participants respond most quickly to. Once all responses have been made, a bias score can be calculated for each participant and then an overall bias can be calculated for the entire sample. If participants are generally faster in their responses for the congruent condition in the obesity IAT example above, this would indicate a negative bias towards obesity related concepts. In analysing response times, the IAT scoring algorithm calculates the difference in average response latency between the congruent and incongruent conditions and dividing by the standard deviation of all latencies for both conditions (Nosek, Greenwald & Banaji, 2007). For paper-based versions of the IAT the difference is calculated based on the number of correct responses in a 20 second period in the congruent/incongruent conditions (e.g., O'Brien et al, 2007). If there is a large difference between the categorisation conditions in terms of response times or number of correct responses, this will result in a larger bias score. The difference between congruent and incongruent response times or accuracy reflects the extent to which people believe that fat people are stupid and thin people are smart. The IAT is thus assumed to offer a window into deeply-rooted beliefs and prejudices that are otherwise difficult to impossible to access explicitly.

The IAT has been used to uncover and measure biases in a wide range of domains, such as attitudes towards race (Greenwald, McGee & Schwartz, 1998), gender stereotyping (Rudman & Kilianski, 2000), alcohol (Wiers et al, 2011) and doping among athletes (Brand et al, 2011) to name but a few examples. What's more, the IAT has been shown to be predictive of people's overt behaviors, underscoring its practical utility. In a meta-analysis of 156 studies, Greenwald and colleagues (Greenwald, Poehlman, Uhlmann, & Banaji, 2009) found that IAT measures correlated significantly with explicit measures of behaviour. In some cases, IAT scores are better predictors of behaviour than more explicit measures. For example, Asendorpf, Banse, and Mücke (2002) found that shyness in individuals was better predicted by a shyness-oriented IAT than by explicit self-ratings of shyness. In some ways the IAT seems to capture attitudes and beliefs that we hold, but that which find it difficult to consciously and explicitly access ourselves.

The Current Study

We have described the IAT paradigm and suggested how it may draw on processing from either the shallow linguistic system or be reliant on deeper processing from the simulation system. To examine whether IAT performance is predicted by linguistic associations we used behavioural data from several published IAT studies and linguistic data extracted from the World Wide Web, using the Web 1T 5-gram corpus (Brants & Franz, 2006). The Web 1T is a snapshot of web pages indexed by Google in 2006 and contains over 1 trillion tokens, making it one of the most representative corpora of language available. Our aim was to examine whether co-occurrence patterns in the linguistic data could predict the effect sizes observed in the behavioural data. We expected that if implicit attitudes (as captured by the IAT) reflect the distributional patterns of language in the linguistic system, then we should see a significant fit between the two sets of data using regression analyses. On the other hand, if the IAT relies on deeper processing in the simulation system, we would not expect such a relationship to exist.


IAT Topic	Reference	N
Race (2)	Rudman & Ashmore, 2007	128
Flowers vs Insects	Greenwald et al., 1998	32
Instrument vs Weapon	Greenwald et al., 1998	32
Japanese vs Korean (2)	Greenwald et al., 1998	64
Alcohol	Wiers et al., 2011	108
Doping (2)	Brand et al., 2011	102
Alcohol and Sport	O'Brien & Lynott, in prep	120
Obesity (3)	O'Brien et al, 2007	1032
Gender (3)	Rudman & Kilianski, 2000	207

Table 2: Sources of the 16 Implicit Association Tests used in the study, with number of studies from each source in brackets, and the total N participants for each source.

Method

Materials We selected 16 IATs from journal articles that used lexical (as opposed to pictorial) stimuli and provided a full list of materials used. Table 2 lists the topics of investigation for each of the IATs, the source reference and the total sample sizes from the original studies; 15 of the studies come from published journal articles and one from a paper in preparation (O'Brien & Lynott, in prep). The IATs cover a broad range of topic areas, including studies of racial stereotypes, obesity and gender roles. Each IAT consists of a set of category concepts (e.g., male, female, good, bad) and a set of target attributes that are of positive or negative valence (e.g., reliable, pleasant, terrible, nasty). The number of attributes used by the IATs ranged from 3 to 25, with a total of 324 target stimuli.

Figure 1: Scatterplot, with line of best fit, for IAT scores from human behavioural data plotted against scores derived from log-transformed linguistic data.



Human Behavioural Data The 16 IATs represent data from 1825 participants. From each IAT we extracted the overall effect size of the bias found based on the human responses (D_H). The IAT effect size is closely related to Cohen's d , a popular measure of statistical effect size.

Linguistic Model The aim of the model is to calculate the size of the linguistic bias (D_L) for each IAT, based on the specific terms used in each task. In order to approximate the linguistic distributional information, we carried out a corpus analysis using the Web 1T 5-gram corpus. Using the corpus we are able to calculate the difference in the strength of associations for each categorisation condition in each IAT. For example, for an IAT examining preference for flowers versus insects, we calculate the level of association between one categorisation pairing (e.g., co-occurrences for "flower" and all positive attributes + co-occurrences for "insect" and all negative attributes) and the inverse, "incongruent" categorisation pairing (co-occurrences for "flower" and all negative attributes + co-occurrences for "insect" and all positive attributes).

The strength of association is calculated by summing the frequency of co-occurrences of category terms (e.g., flower, insect) with each of the attribute terms (e.g., nice, horrible, etc.). For each category-attribute pairing (e.g., flower-nice), we calculated the cumulative 5-gram frequency of forward and backward co-occurrences between the category word and attribute (i.e., the summed count of occurrences of

[flower ... nice] and [nice ... flower] in the corpus with zero, one, two and three intervening words: for a similar approach, see Louwerse & Connell, 2011). Because of the large number of calculations this required (>10,000), we developed a semi-automated tool to take each set of IAT terms and output the collocation frequencies from the Web 1T corpus. Using these summed frequencies we then calculated a linguistic effect size using two models; one model based on raw frequency counts and one model based on the log-transformed frequencies (using the natural log). Ji (2010) discusses the improvement in the distribution curves of datasets of seven sub-corpora after having undergone log transformation as the transformation mitigates the effects of extreme values (i.e., highly frequent terms). Thus, both linguistic models represent the ratio of the frequencies for the congruent categorisation condition compared to the frequencies of the incongruent categorisation condition. Finally, we conducted separate linear regression analyses for the two linguistic models using the linguistic bias (D_L) as a predictor variable and the behavioural bias (D_H) as the dependent variable.

Results & Discussion

The effect sizes in the human data ranged from -.49 to 3.65 ($M = 1.2$, $SD = 1.07$), while the effect sizes in the linguistic data ranged from .2 to 7.3 ($M = 2.02$, $SD = 2.1$) for the raw frequency model and from .85 to 1.31 ($M = 1.05$, $SD = .14$)

for the log-transformed model. Using linear regression analyses we found significant relationships between the effect sizes calculated from the human data, D_H , and the effect sizes calculated from the linguistic models, D_L . This positive relationship indicates that the larger the effect in the linguistic data, the larger the predicted effect in the human data. Figure 1 illustrates the relationship between the biases predicted from the log ratio linguistic model and those derived from the human data. The regression model for the raw frequency ratio model was significant ($r^2 = .612, p < .001, n = 16$) resulting in a β -coefficient of .782 for the linguistic predictor ($t = 4.696, p < .001$). The regression model for the log frequency ratio model was also significant ($r^2 = .596, p < .001, n = 16$) resulting in a β -coefficient of .772 for the linguistic predictor ($t = 4.544, p < .001$). This indicates that both models reflect approximately 60% of the variance in the human IAT scores.

General Discussion

This study investigated whether linguistic distributional information can be used to predict levels of implicit attitudes as measured by the Implicit Association Test. We observed significant relationships between effect sizes from human behavioural data and effect sizes calculated from linguistic distributional data. This finding suggests that performance on the IAT may not reflect the deeply-rooted biases and beliefs held by individuals and groups, but instead reflects shallower linguistic associations they have encountered in their environment. While the present results are promising, there are of course some caveats we need to be aware of. We discuss some of these limitations and avenues for future research below.

An obvious question to ask is, if the IAT reflects only shallow linguistic processing, then how is IAT performance predicting overt behaviours? There are two issues here. The first is that the while IAT is successful in predicting outcomes in certain sub-domains (e.g., political preferences), it poorly reflects outcomes in others (e.g., sexual orientation; Greenwald et al., 2009). The second is that even where IAT performance is claimed to predict other behavioural outcomes, the claims may not stand up to closer scrutiny. It is important when comparing overt measures to IAT performance that one uses implicit tasks that have clear behavioural outcomes. Good examples of this are patient treatments and using realistic CVs/resumés for assessing job candidates. Green et al (2007) found that doctors' levels of implicit bias towards black patients did not always tally with their decision to offer treatment using thrombolysis to remove blood clots. Although doctors with higher anti-black bias were less likely to treat black patients than white patients using thrombolysis, those doctors with low anti-black bias (but not a *pro*-black bias) were actually more likely to treat black patients than white patients. In re-analysing data looking at racially discriminating behaviour in job candidate selection, Blanton and colleagues (2009) found that the IAT failed to predict any discriminatory behaviour when factors such as rater reliability and outlier removal were taken into account. In one case, previous

evidence of anti-black prejudice was actually reversed revealing a pattern of pro-black bias. However, the process of candidate evaluations can be broken down further. For example, selection and shortlisting of candidates can be viewed as more of a heuristic process, while providing specific grades to individual candidates can be seen as more deliberative. Blommaert, van Tubergen and Coenders (2011) distinguished between these two aspects of the candidate assessment process while examining the effects of implicit attitudes towards ethnicity and gender. They found that only explicit measures predicted people's grading of candidates, but that both implicit (IAT) and explicit measures predicted people's shortlisting of candidates. Thus, it is important to take into account both the task domain and the nature of the task (i.e., heuristic or deliberative) to have a clearer idea of whether the IAT and therefore linguistic distributional information may have a role to play in predicting behavioural outcomes.

A limitation of the current approach is that it does not discriminate between different groups and different task contexts and how this would affect performance on a given IAT. For example, we would not necessarily expect a group of American students and a group of Chinese students to show the same type of bias judging American and Chinese faces paired with positive or negative attributes. One possibility would be to extend the model to incorporate additional domain terms that would calculate co-occurrence frequencies but limited to specific contexts to attempt to approximate these contextual effects.

Although it may be argued that language reflects our cultural beliefs and social norms, it is difficult to establish a causal relationship between implicit attitudes and the linguistic data. It may be that the linguistic distributional information is a) driving the formation of these biases, b) is the behavioural outcome of these biases or c) part of a self-sustaining cycle of biases influencing language influencing biases and so on. However, it is clear that exposure to socio-cultural attitudes does impact our own attitudes and behaviours, as developmental changes are evident in IAT performance. For example, older subjects tend to show larger IAT effects than younger subjects (Hummert, Garstka, O'Brien, Greenwald, & Mellott, 2002). As language is one of the key methods for the transmission of socio-cultural information, this underscores the possible role for language exposure in the formation of implicit attitudes (Nosek, Greenwald & Banaji, 2007).

In conclusion, we present the first model of implicit attitudes based on linguistic data extracted from the world wide web. We found that linguistic models revealed a strong correspondence with human behavioural data. We see language as a primary means for transmission of attitudinal information and agree with Uhlman and colleagues (in press) that "implicit attitudes reveal the power of cultures to reproduce themselves in individual minds". However, our findings also suggests that such implicit attitudes may not represent deeply rooted beliefs as has previously been assumed. Ongoing work is exploring the predictive power of the model by using linguistic data to predict attitude

effect sizes in advance of behavioural studies, providing a strong test whether our hidden beliefs can be revealed in the patterns of our language use.

References

- Asendorpf, J. B., Banse, R., & Mücke, D. (2002). Double dissociation between implicit and explicit personality self-concept: The case of shy behavior. *Journal of Personality and Social Psychology*, *83*, 380–393.
- Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. De Vega, A. M. Glenberg, & A. C. Graesser, A. (Eds.). *Symbols, embodiment, and meaning*. Oxford, UK: OUP.
- Blanton, H., Jaccard, J., Klick, J., Mellers, B., Mitchell, G., & Tetlock, P. E. (2009). Strong Claims and Weak Evidence: Reassessing the Predictive Validity of the IAT. *Journal of Applied Psychology*, *94*, 567–582.
- Blommaert, L., van Tubergen, F., & Coenders, M. (2011). Implicit and explicit interethnic attitudes and ethnic discrimination in hiring. *Social Science Research*, *41*, 61–73.
- Brand, R., Melzer, M. & Hagemann, N. (2011). Towards an implicit association test (IAT) for measuring doping attitudes in sports. Data-based recommendations developed from two recently published tests. *Psychology of Sport and Exercise*, *12*, 250 – 256.
- Brants, T., & Franz, A. (2006). *Web 1T 5-gram Version 1*. Philadelphia: Linguistic Data Consortium.
- Connell, L., & Lynott, D. (2011). Modality switching costs emerge in concept creation as well as retrieval. *Cognitive Science*, *35*, 763-778.
- Connell, L., & Lynott, D. (2012). Principles of Representation: Why You Can't Represent the Same Concept Twice. Under Review.
- Fazio, R. H. & Olsen, M. A. (2003). Implicit measures in social cognition research: Their Meaning and Use. *Annual Reviews in Psychology*, *54*, 297–327
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013–1027
- Foroni, F. & Semin, G. R. (2012). Not all implicit measures of attitudes are created equal: Evidence from an embodiment perspective. *Journal of Experimental Social Psychology*, *48*, 424–427
- Furnham, A. (1986). Response bias, social desirability and dissimulation. *Personality and Individual Differences*, *7*, 385-400.
- Green, A. R., Carney, D. R., Pallin, D. J., Ngo, L. H., Raymond, K. L., ... (2007). Implicit Bias among Physicians and its Prediction of Thrombolysis Decisions for Black and White Patients. *Society of General Internal Medicine*, *22*, 1231–1238.
- 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., Poehlman, T.A., Uhlmann, E.L., & 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.
- Hummert, M. L., Garstka, T. A., O'Brien, L. T., Greenwald, A. G., & Mellott, D. S. (2002). Using the implicit association test to measure age differences in implicit social cognitions. *Psychology and Aging*, *17*, 482–495.
- Ji, M. (2010). A corpus-based study of lexical periodization in historical Chinese. *Literary and Linguistic Computing*, *25*, 199 – 213.
- Louwerse, M. M., & Connell, L. (2011). A taste of words: Linguistic context and perceptual simulation predict the modality of words. *Cognitive Science*, *35*, 381-398.
- Louwerse, M. M., & Jeuniaux, P. (2008). Language comprehension is both embodied and symbolic. In M. de Vega, A. Glenberg, & A. C. Graesser (Eds.), *Symbols, embodiment, and meaning*. OUP.
- Louwerse, M.M. & Jeuniaux, P. (2010). The linguistic and embodied nature of conceptual processing. *Cognition*, *114*, 96-104.
- Lynott, D., & Connell, L. (2010). Embodied conceptual combination. *Frontiers in Psychology*, *1*:216, 1-14.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Harvesting Implicit Group Attitudes and Beliefs From a Demonstration Web Site. *Group Dynamics: Theory, Research, and Practice*, *6*, 101-115.
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). The Implicit Association Test at age 7: A methodological and conceptual review. In J. A. Bargh (Ed.), *Social Psychology and the Unconscious: The Automaticity of Higher Mental Processes* (pp. 265-292). New York: Psychology Press.
- O'Brien, K. S., Hunter, J. A., & Banks, M. (2007). Implicit anti-fat bias in physical educators: Physical attributes, ideology, and socialisation. *International Journal of Obesity*, *31*, 308– 314.
- O'Brien, K. S., Hunter, J. A., Halberstadt J., & Anderson J. (2007). Body image and explicit and implicit anti-fat attitudes: the mediating role of physical appearance comparisons. *Body Image*, *4*, 249–256.
- Paulhus, D. L. (1991). Measurement and control of response bias. In: Robinson JP, Shaver PR, Wrightsman LS (eds). *Measures of Personality and Social Psychological Attitudes*. Academic Press: New York, 1991, pp 17–59.
- Rudman, L.A. & Kilianski, S.E. (2000). Implicit and Explicit Attitudes Toward Female Authority. *Personality and Social Psychology Bulletin*, *26*, 1315 – 1328.
- Solomon, K.O., & Barsalou, L.W. (2004). Perceptual simulation in property verification. *Memory & Cognition*, *32*, 244-259.
- Uhlmann, E.L., Poehlman, T.A., & Nosek, B. A. (in press). Automatic Associations: Personal Attitudes or Cultural Knowledge? To appear in Jon D. Hanson (Ed.) *Ideology, Psychology, and Law*. New York, NY: OUP.
- Wiers, R.W., Eberl, C., Rinck, M., Becker, E.S. & Linenmeyer, J. (2011). Retraining Automatic Action Tendencies Changes Alcoholic Patients' Approach Bias for Alcohol and Improves Treatment Outcome. *Psychological Science*, *22*, 490 – 497.