But for the Bad, There Would Not Be Good: Grounding Valence in Brightness Through Shared Relational Structures

Daniël Lakens
Eindhoven University of Technology

Gün R. Semin
Utrecht University and Koç University

Francesco Foroni
Utrecht University

Light and dark are used pervasively to represent positive and negative concepts. Recent studies suggest that black and white stimuli are automatically associated with negativity and positivity. However, structural factors in experimental designs, such as the shared opposition in the valence (good vs. bad) and brightness (light vs. dark) dimensions might play an important role in the valence–brightness association. In 6 experiments, we show that while black ideographs are consistently judged to represent negative words, white ideographs represent positivity only when the negativity of black is coactivated. The positivity of white emerged only when brightness and valence were manipulated within participants (but not between participants) or when the negativity of black was perceptually activated by presenting positive and white stimuli against a black (vs. gray) background. These findings add to an emerging literature on how structural overlap between dimensions creates associations and highlight the inherently contextualized construction of meaning structures.

Keywords: valence, brightness, knowledge structures, emergent meaning, context effects

Supplemental materials: http://dx.doi.org/10.1037/a0026468.supp

In Plato's chariot allegory of the journey of the soul, the charioteer (or reason) has to control two horses: One is white, representing the moral impulse, and the other is black, representing the immoral passions. The opposition between light and dark is used pervasively to metaphorically structure the opposition between the abstract concepts of good and bad. Recently, researchers have proposed that the brightness of stimuli can lead to automatic inferences about their valence, with bright stimuli activating associations with positivity and dark stimuli activating associations with negativity (Meier, Robinson, & Clore, 2004; Sherman & Clore, 2009). At the same time, the literature on color evaluations suggests that white is affectively neutral. Given this discrepancy, we examine the alternative hypothesis that instead of an automatic association between white and positivity, the positivity of white emerges when the shared relational structures of the light–dark and positive–negative bipolar oppositions are coactivated in an experimental paradigm. The positivity of white is predicted to emerge only in opposition to the negativity of black.

In this article, we propose that inconsistencies in the evaluation of white can be resolved with a shared relational structures view that details how mappings between stimuli and responses emerge within experimental tasks due to alignment processes that construct relationships between endpoints of bipolar dimensions that are manipulated within participants. In line with recent examples of how structural factors in experimental designs can lead to mappings between stimuli and responses (e.g., Duscherer, Holender, & Molenaar, 2008; Lakens, 2011; Proctor & Cho, 2006; Rothermund & Wentura, 2004; Scherer & Lambert, 2009), the shared polar oppositions in the valence and brightness dimensions might activate negative–black and positive–white relationships not due to intrinsic associations but because the opposition in the brightness dimension is used to structure the opposition in the valence dimension. By manipulating valence and brightness either within or between participants, the current series of experiments highlights how structural factors in experimental designs can allow mappings between stimuli and responses to emerge or prevent them from doing so.

Explicit Brightness–Valence Associations

Color research suggests a clear anchoring of black as negative. For example, referees judge sports teams wearing black outfits as more aggressive and consequently penalize these teams more often (Frank & Gilovich, 1988). The negativity of black has been argued to result from the fact that humans are diurnal animals (Williams & Morland, 1976). This evolutionary explanation is supported by

research showing that people have a general tendency to feel less at ease in the dark, as signified by an increase in the human startle reflex in the dark compared to normal lighting conditions (Grillon, Pellowski, Merikangas, & Davis, 1997). Perhaps surprisingly, color researchers have found that white is generally judged to be affectively neutral (e.g., Burkitt, Barrett, & Davis, 2003; Götz & Götz, 1974; McManus, Jones, & Cottrell, 1981).

The notion that white is associated with positivity does not originate from color research but from cross-cultural studies on the measurement of meaning, which uses a method known as the semantic differential (e.g., Adams & Osgood, 1973; Osgood, 1960; Osgood, Suci, & Tannenbaum, 1957). Participants in these studies are asked to indicate whether a given stimulus is most clearly related to one of two bipolar opposites (e.g., large–small, bright–dark). Importantly, this measurement procedure defines the meaning of concepts as a location in a multidimensional semantic space consisting of pairs of bipolar adjectives (and is therefore also referred to as the meaning-polarity test; Osgood, 1952). Consequently, white and black are functionally treated as opposites in these measurements (Osgood et al., 1957), and it is within this field of oppositions that positivity is represented by brightness but not darkness. This interpretation is also supported by the results of three pilot studies we performed with the aim of providing further empirical support for the neutral evaluation of white and the negative evaluation of black in our participant population (see the supplementary online material). These studies clearly confirm that black is evaluated negatively, whereas white is evaluated as a neutral color.

To conclude: There is a strong association between darkness and negativity, irrespective of the context. However, white is only a representation of positivity when evaluations are made within a context where the positive–negative and light–dark oppositions are coactivated (i.e., when the shared bipolar oppositions in the brightness and valence dimensions are present and salient).

One-to-One Grounding

Considerable research over the past few years has revealed how concepts related to affective experiences are structured in concrete dimensions such as space and brightness (for a review, see Crawford, 2009). For example, Meier and colleagues (2004) investigated the automatic nature of the valence–brightness metaphor by examining whether task-irrelevant perceptual characteristics of stimuli (i.e., their brightness) influenced speeded evaluations of stimulus words. The authors predicted that positive words in white and negative words in black (so-called metaphor congruent mappings) would be categorized faster and more accurately than positive words written in black and negative words written in white (so-called metaphor incongruent mappings). These predictions were confirmed, and Meier and colleagues concluded that their results point to "the obligatory nature of affective inferences based on stimulus brightness" (Meier et al., 2004, p. 85).

Metaphor congruency effects are often explained as the result of Stroop-like interference processes (see Meier & Robinson, 2004, Meier et al., 2004; Schubert, 2005; Sherman & Clore, 2009). Due to the theorized obligatory nature of associations between perceptual characteristics of stimuli (e.g., their brightness) and their conceptual meaning (e.g., their valence), brightness is hypothesized to lead to automatic inferences concerning the valence of the

stimulus (e.g., positive = white, negative = black). During a speeded evaluation task, the automatic inference of positivity when presented with a bright stimulus should facilitate positive evaluations but interfere with negative evaluations. As Meier and colleagues (2004, p. 86) explained, "a negative word presented in white would give rise to two response tendencies, one to respond 'negative' (on the basis of stimulus valence) and one to respond 'positive' (on the basis of stimulus color)."

The conclusions previous researchers have drawn from their findings are based on the theoretical perspective that relationships between concrete experiences (e.g., brightness) and abstract concepts (e.g., valence) are fixed neural connections representing strong and stable associations between concepts and experiences (Lakoff, 2008). The translation of perceptual information into conceptual meaning is argued to be "used constantly and automatically, with neither effort nor awareness" (Lakoff, 1993, pp. 227–228), and is assumed to be obligatory at the stage of word encoding (Meier & Robinson, 2004). Thus, perceiving the color white should be sufficient to activate its associated valence, independent of the context in which it is encountered. We call this theoretical perspective *one-to-one grounding* because the prevalent assumption in these investigations is that the observed effects reflect stable and direct associations in long-term memory.

Shared Relational Structures

An alternative approach to the metaphoric mapping of valence and brightness presupposes that the main function of metaphors is to "provide relational structure to those domains where the structure may not be obvious from world experience" (Boroditsky, 2000, p. 3). By importing the polar opposition from the concrete brightness dimension to structure the more abstract relationship between positivity and negativity in the valence dimension (see Boroditsky, 2000; Gattis, 2001; Gentner & Bowdle, 2008; G. Murphy, 1996),[1] mappings between endpoints of the valence and brightness dimension can emerge because of their shared relational structures. Following such reasoning, which we refer to as a *shared relational structures view,* white can become associated with positivity, but such a mapping is highly context specific, and it is hypothesized to become activated primarily in opposition to the negativity of black. The negativity of black is assumed to be context independent due to its theorized evolutionary origin (see

---

[1] Although Lakoff and Johnson (1980) mainly discussed metaphors that relate one concept in terms of another concept (e.g., an argument is a building), they noted that some metaphors structure a whole system of concepts (e.g., valence) with respect to another domain (e.g., verticality). These so-called *orientational metaphors* often structure abstract concepts in concrete dimensions, such as the polar oppositions up–down, in–out, or light–dark. However, the relational structure in these metaphors (typically consisting of polar oppositions) did not receive much attention, nor were such metaphors conceptually differentiated from other types of metaphors, and this idea was no longer present in later work (e.g., Lakoff & Johnson, 1999). We therefore conclude that the relational structures in these orientational metaphors are not seen to have a special status beyond the two complementary metaphors (e.g., light = good and dark = bad). Thus, even relational metaphors seem to be best described as an example of one-to-one grounding, instead of metaphorically structuring a shared relational structure (e.g., the relationship between light and dark is similar to the relationship between good and bad).

Scherer & Lambert, 2011, for recent work on the context dependency of white and the context independency of black in implicit race biases). This context specificity of relations between stimuli and evaluations is in line with recent theoretical work stressing that evaluation is a constructive process (Schwarz, 2007), a viewpoint that is central in situated cognition perspectives (e.g., Smith & Semin, 2004).

The one-to-one grounding account and the shared relational structures view differ on the necessary preconditions for the positivity of white to emerge. The one-to-one grounding account is based on the notion that white is associated with positivity and black is associated with negativity irrespective of the context in which these stimuli are encountered, while the shared relational structures view suggests that the positivity of white only emerges in opposition to the negativity of black. Since the studies that examined the automatic association between valence and brightness (e.g., Meier et al., 2004; Sherman & Clore, 2009; but see also Smith-McLallen, Johnson, Dovidio, & Pearson, 2006) relied on designs where the brightness and valence dimensions were always manipulated within participants (e.g., each participant had to categorize both positive and negative words, presented both in a black and in a white font), these studies invariably coactivated the positive–negative and light–dark oppositions. Thus, the question whether black or white stimuli in isolation are associated with negativity or positivity remains unanswered, making it impossible to distinguish between the one-to-one grounding account and the shared relational structures view based on the research conducted to date.

How Polar Oppositions Create Shared Relational Structures

The importance of bipolar oppositions in experimental tasks is not limited to investigations of conceptual metaphors. Indeed, theoretically similar issues have recently emerged as a topic of interest in other experimental paradigms in implicit social cognition. For example, Moors, Spruyt, and De Houwer (2010, p. 29) noted how the mere fact that "all stimuli in a prototypical affective priming study have a polarized valence might be sufficient to increase the salience of valence." In many experimental paradigms, there is not only a polar opposition in the stimulus set (e.g., positive vs. negative words) but also a polar opposition in the response categories (e.g., approach vs. avoidance movements, positive vs. negative evaluations). In the current article, we propose that the oppositions in the stimulus and response dimensions will increase the salience of the presence of shared relational structures between the two dimensions. Where polar opposites in the stimulus dimension (e.g., positive vs. negative) might increase the salience of the valence dimension (cf. Moors et al., 2010), the presence of shared relational structures in the bipolar stimulus and response dimensions might increase the salience of the structural overlap between the two dimensions.

According to a shared relational structures view, structural similarities between stimulus dimensions (e.g., both stimulus dimensions consist of bipolar opposites) can lead to context-specific associations between the endpoints of these dimensions. People implicitly prefer relational structures and are more likely to search for similarities and parallels between dimensions than to assume that two bipolar dimensions in an experimental task are completely orthogonal and unrelated (for a review, see Gentner & Bowdle, 2008). We propose that the shared polar opposition in the conceptual dimension of stimuli (e.g., their positive vs. negative valence) and the perceptual dimension of stimuli (e.g., their white vs. black color) is assumed to lead to an alignment process where white, evaluated neutrally in isolation, is mapped onto positivity in opposition to the negativity of black. Although such structural mappings do not have to be problematic, under specific circumstances, shared relational structures can provide a source of bias in participants' judgments. Such a bias might be easily misinterpreted as a stable attitude represented in memory, while in reality reflecting a temporary mapping due to specific characteristics of the experimental paradigm.

Consider the following three situations. First, black is intrinsically associated with negativity, and white is intrinsically associated with positivity. If this would be the case, the presence of polar oppositions in the response dimensions might increase the strength of the relationship between valence and brightness (see Moors et al., 2010), but at least this mapping would reflect underlying intrinsic associations. A second possibility is that black is not associated with negativity and white is not associated with positivity. Again, the presence of shared relational structures due to the polar oppositions will lead to an alignment process between the valence and brightness dimensions. Since neither black nor white is strongly associated with valence, random relationships between valence and brightness will form, with approximately half of the participants mapping white onto positivity and half of the participants mapping white onto negativity. Overall, the mappings due to shared relational structures will not lead to an observable bias. Now, consider a third, more problematic possibility. Black is strongly associated with negativity, but white is affectively neutral (note that we do not discuss the theoretically equivalent fourth possibility that white is associated with positivity but black is affectively neutral). In such a situation, the shared relational structures in the valence and brightness dimensions will lead to a mapping between black and negativity (reflecting intrinsic associations) and a mapping between white and positivity (purely due to the shared relational structures). In this last case, there is a risk of misinterpreting both of the complementary mappings (positivity–white and negativity–black) in support of a stable association between valence and brightness while only the latter is present. We investigated the possibility that whereas black might be associated with negativity regardless of task characteristics and the configurations of the stimulus set, the mapping between white and positivity emerges purely through shared relational structures.

The Current Research

According to the shared relational structure view, whether or not the positivity of white emerges hinges upon the opposition between the endpoints of the valence and brightness dimensions. Following this rationale, whenever the polar opposition in the conceptual (i.e., valence) or perceptual (i.e., brightness) dimension is removed from the experimental paradigm (e.g., by manipulating valence and/or brightness between instead of within participants) white should no longer be associated with positivity. According to the one-to-one grounding account,

white should always activate associations with positivity, and black should always activate associations with negativity. To compare predictions from a one-to-one grounding account with the shared relational structures view, we manipulated valence and brightness either between or within participants. We developed a novel Chinese translation paradigm where participants were asked to indicate whether Chinese ideographs were the correct translation for a word presented on the screen. The idea of this paradigm is that irrelevant perceptual characteristics (i.e., the brightness of the ideograph) will influence translation judgments for positive and negative words.

As in conceptually related projective tests such as the affect misattribution paradigm (AMP; Payne, Cheng, Govorun, & Stewart, 2005), the Chinese translation paradigm examines whether Chinese ideographs are judged to reflect the affective meaning of the stimuli (see also Cacioppo, Priester, & Berntson, 1993; S. T. Murphy & Zajonc, 1993), but instead of using a uniform group of Chinese ideographs as stimuli as in earlier studies, the perceptual characteristics of the ideographs are manipulated within or between individuals (see also Lakens, Semin, & Foroni, 2011). Projective tasks that rely on judgments under uncertainty (such as the AMP) have proved to be extremely useful for investigating the structural factors that determine responses in experimental tasks (e.g., Deutsch & Gawronski, 2009; Gawronski, Cunningham, LeBel, & Deutsch, 2010). Although a judgment under uncertainty paradigm does not allow us to test whether the brightness–valence metaphor is automatic in nature, our main interest lies in determining whether the valence–brightness metaphor is conditional upon the presence of the shared relational structure between the opposition in the brightness dimension and the opposition in the valence dimension.

Overview of the Studies

In Experiment 1, participants were presented with positive and negative words and black and white Chinese ideographs (manipulating both endpoints of the valence and brightness dimensions within participants). Under these conditions, both a one-to-one grounding account and a shared relational structures view would predict that positive words should be translated above guessing average by white ideographs and that negative words should be translated above guessing average by black ideographs. This outcome would replicate the pattern observed in earlier research (Meier et al., 2004) with a new paradigm. In Experiments 2–4, the two accounts were pitted against each other. In Experiment 2, brightness was manipulated within participants, but valence was manipulated between participants. In Experiment 3, brightness was manipulated between participants while valence was manipulated within participants, and in Experiment 4, both brightness and valence were manipulated between participants. The one-to-one grounding account of valence in brightness would predict similar results across Experiments 2–4, reproducing the pattern obtained in Experiment 1. On the other hand, the shared relational structures view predicts that the probability of translating positive words by white ideographs will not differ from guessing average because the polar opposition in the perceptual and conceptual dimensions is not activated within subjects. Thus, the key comparison

between the two accounts is based on the prediction that the positivity of white should not emerge across Experiments 2–4.

The salience of bipolar oppositions in stimulus dimensions can be influenced by manipulating factors within or between participants but also by adding additional neutral trials to the stimulus set when manipulating valence and brightness within participants (for related results in the Stroop task, see Kahneman & Chajczyk, 1983), which replaces the binary opposition in the valence dimension with an ordinal structure (positive, neutral, negative). We predicted that by adapting Experiment 1 to include neutral stimuli, structural mappings would no longer emerge, reducing the likelihood that participants would choose a white ideograph to translate positive words. Experiment 6 aimed to provide direct support for the hypothesis that the positivity of white emerges in opposition to the negativity of black. Participants were presented with positive words and white ideographs, either against a gray background (identical to Experiment 4) or against a (negatively evaluated) black background. The polar oppositions in the valence and brightness dimensions were present in the black background condition but not in the gray background condition. We predicted that positive words would be translated above guessing average by white ideographs but only in the black background condition. The hypothesis in Experiment 6 follows only from the shared relational structures view and cannot be explained by a one-to-one grounding account.

Experiment 1

To establish a paradigmatic baseline in the Chinese translation paradigm, we aimed to replicate the pattern of results observed by Meier et al. (2004) by manipulating valence (positive vs. negative) and brightness (white vs. black) in a within-participants design. Positive words were expected to be translated above guessing average by white ideographs, and negative words were expected to be translated above guessing average by black ideographs.

Method

Participants. Thirty-four paid voluntary students (23 females, mean age 20 years) at a Dutch university took part in this experiment, which conformed to a 2 (word valence: positive vs. negative) × 2 (ideograph brightness: white vs. black) within-participants design, and received €3 ($4.50) for their participation.

Procedure. Participants were seated in individual cubicles. All instructions were presented on a computer screen in dark-gray letters on a perfectly gray background (red–green–blue: 128, 128, 128). Participants were informed that recent studies had revealed that people who did not speak Hebrew nevertheless had a 60% chance of correctly translating Hebrew words. The current experiment would ostensibly investigate if people were able to score above guessing average when translating Dutch words to Chinese ideographs. Before starting the translation task, a Chinese ideograph (not used in the translation task) was presented to the participants: If they indicated they could read Chinese ideographs, they were removed from the analysis. This general instruction and procedure were identical across all six experiments presented here.

Participants received 12 stimulus words (six positive, six negative)[2] presented in random order. Underneath each word, one white and one black Chinese ideograph were displayed. Participants were asked to judge which of the two ideographs correctly represented the meaning of the Dutch stimulus word. Twenty-four ideographs were randomly displayed in either white or black.[3] There was no time limit for the translation judgments. After performing the 12 translation judgments, participants were asked to rate the valence of the 12 words on a scale from 1 (*very negative*) to 7 (*very positive*).

Results

Manipulation check. Positive words were judged as more positive ($M = 6.51$) than negative words ($M = 1.80$), and a paired-samples t test indicated this difference was significant, $t(32) = 29.06$, $p < .001$.

Translation judgments. One participant was excluded from the analysis because she spoke Chinese. Since the choices for white or black ideographs were mutually dependent, we only calculated the average number of times positive and negative words were translated by a white ideograph. As expected, participants translated the six positive words above guessing average by white ideographs (61.50%, $M = 3.69$, $SD = 1.24$), as indicated by a test against chance, $t(32) = 3.24$, $p = .003$, Cohen's $d = .57$. The six negative words were translated below guessing average by white ideographs, (40.34%, $M = 2.42$, $SD = 1.09$), $t(32) = -3.03$, $p = .005$, Cohen's $d = .54$. A repeated measures analysis of variance (ANOVA) with word valence as the only within-subject factor indicated participants translated positive words more often by a white ideograph compared to negative words, $F(1, 32) = 15.47$, $p < .001$, $\eta_p^2 = .33$ (see Figure 1, top panel).

Discussion

Even though the brightness of the ideographs was irrelevant for the translation task, differences in the brightness of the ideographs influenced the likelihood with which they were judged to correctly translate positive or negative words. As such, this experiment conceptually replicated previous findings (Meier et al., 2004; Sherman & Clore, 2009) in a paradigm based on judgments under uncertainty.

Experiments 2–4 were performed to investigate whether the bias to translate positive words by white ideographs would still be observed when the polar oppositions in the valence and brightness dimensions were removed from the translation task by manipulating valence, brightness, or both dimensions in a between- instead of within-participants design.

Experiment 2

In Experiment 2, the valence dimension was manipulated between participants. If the shared relational structure in Experiment 1 contributed to the observed bias in the translation judgments of participants, breaking up the positive–negative opposition in the current experiment should influence the translation judgments. Although negative words were still expected to be translated above guessing average by black ideographs, positive words should no longer be translated above guessing average by white ideographs.
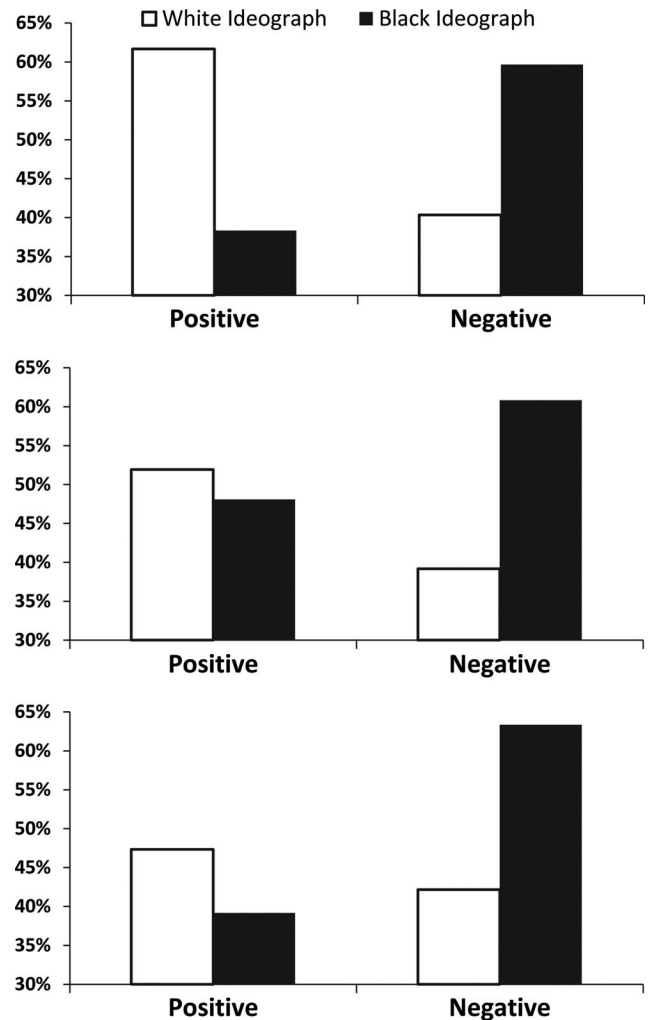

Figure 1. Percentages of choices for black versus white ideographs (mutually exclusive in Experiment 1, top panel, and Experiment 2, middle panel) or *yes* judgments (Experiment 3, bottom panel) compared to guessing average (50%).

[2] Positive and negative stimuli used in the six studies: In Experiments 1, 3, 5, and 6, only words marked with an asterisk were used. Positive words: Baby*, Brave*, Hero*, Love*, Sincere*, Loyal*, Party, Humour, Friend, Nice, Shining, Pleasurable. Negative words: Death*, Cruel*, Poison*, Sick*, Enemy*, Hate*, Criminal, Cancer, Disgust, Unfair, Mean, Filthy. Neutral words: Identical*, Table*, Sandal*, Legal*, Even*, Chair*, Solid, Rug, Glove, Balcony, Portable, Audible.

[3] In a pilot study, 65 students rated the valence of 10 black or white Chinese ideographs on a 7-point scale from −3 (*very negative*) through 0 (*neutral*) to 3 (*very positive*), with the brightness of the ideographs manipulated between participants. The average rating was calculated for the ideographs, and a two-sample t test was used to compare the differences between ratings of the black and white ideographs. The valence ratings for the two ideographs did not differ between white ($M = 0.06$, $SD = 0.48$) and black ($M = 0.09$, $SD = 0.57$) ideographs, $t(63) = 0.17$, $p = .87$, *ns*. Although black and white ideographs do not differ in their explicit evaluations, this does not imply they are equally likely to be chosen to represent the meaning of positive or negative words in a projective task such as the Chinese translation paradigm—it simply means they have no strong evaluative connotation on their own.

If, on the other hand, white ideographs are associated with positivity irrespective of the presence of shared relational structures, manipulating valence between participants should have no influence on the translation judgments, and a pattern of results identical to that in Experiment 1 should be observed.

Method

Participants and procedure. Fifty-three university students (31 females, mean age 20 years) took part in this experiment for monetary compensation (€3, or $4.50) and were randomly assigned to either the positive- or the negative-word condition. The procedure was identical to Experiment 1, with the exception that valence of the stimulus words was manipulated between participants, such that participants were presented with 12 positive or 12 negative stimulus words.

Results

Manipulation check. Positive words were judged as more positive ($M = 6.33$) than negative words ($M = 2.08$), $t(51) = 27.48$, $p < .001$.

Translation judgments. The average number of times participants in the negative-word condition chose the white ideograph to translate the stimulus word was significantly lower than guessing average, (39.17%, $M = 4.70$, $SD = 2.84$), $t(26) = -2.37$, $p = .03$, Cohen's $d = .46$, indicating they preferred the black ideograph to translate negative words. As expected from a shared relational structures view, in the positive-word condition the likelihood with which participants translated positive words by white ideographs did not differ from guessing average, (51.92%, $M = 6.23$, $SD = 1.88$), $t(25) = 0.63$, $p = .54$, Cohen's $d = .01$. A one-way ANOVA revealed that participants in the negative-word condition were less likely to translate the words by white ideographs compared to the participants in the positive-word condition, $F(1, 51) = 5.28$, $p = .026$, $\eta_p^2 = .09$ (see Figure 1, middle panel).

Discussion

If the valence dimension is manipulated between participants, positive words are no longer translated by white ideographs above guessing average, whereas negative words are still translated by black ideographs above chance level. These results are in line with the shared relational structures view according to which the presence of the shared opposition in the valence and brightness dimensions is an essential prerequisite for the positivity of white to emerge.

Experiment 3

Whereas valence was manipulated between participants in Experiment 2, the current experiment compared the one-to-one grounding account with the shared relational structures view by manipulating brightness between participants. The translation task was adapted to enable the manipulation of ideograph brightness between participants. Instead of choosing between a black and a white Chinese ideograph, only one ideograph was displayed beneath each stimulus word, and participants were asked to judge whether the ideograph was a correct translation of the stimulus word by clicking the *yes* or *no* button. On the basis of the shared

relational structures view, we predicted that participants would judge that black ideographs correctly translated negative words above chance level. On the basis of the fact that the opposition in the brightness dimension was not present in the translation task, no deviations from guessing average were expected in the white-ideograph condition.

Method

Participants and procedure. Thirty-nine students (30 females, mean age 21 years) at a Dutch university took part in this experiment for monetary compensation (€3, or $4.50) and were randomly assigned to the white- or black-ideograph condition. Participants were asked to indicate if they thought the displayed Chinese ideograph was the correct translation for the stimulus word (or not) by clicking on the *yes* (or the *no*) buttons to the left and the right of the ideograph. The location of the *yes* and *no* buttons on the screen was randomized in each of the 12 trials. The same stimulus words were used as in Experiment 1.

Results

We calculated the average number of *yes* responses for positive and negative words in the black- and white-ideograph conditions. Since responses were no longer mutually exclusive, as in Experiment 1 and 2, where choosing the white ideograph meant participants did not choose the black ideograph, all four averages were tested against chance. In the white-ideograph condition, participants' average number of *yes* judgments did not differ from guessing average for positive (47.34%, $M = 2.84$, $SD = 1.74$) or negative (42.17%, $M = 2.53$, $SD = 1.39$) stimulus words, $t(18) = -0.40$, $p = .70$, Cohen's $d = .09$, and $t(18) = -1.49$, $p = .16$, Cohen's $d = .34$, respectively. In the black-ideograph condition, participants translated positive words by black ideographs below guessing average (39.17%, $M = 2.35$, $SD = 0.99$), $t(19) = -2.94$, $p = .008$, Cohen's $d = .66$, and negative words above guessing average (63.34%, $M = 3.80$, $SD = 1.24$), $t(19) = 2.89$, $p = .009$, Cohen's $d = .65$ (see Figure 1, bottom panel).

A repeated measures analysis with ideograph color (white vs. black) as a between-participant factor and word valence (positive vs. negative) as a within-participant factor revealed only a significant interaction of word valence and stimulus brightness, $F(1, 37) = 8.58$, $p = .006$, $\eta_p^2 = .18$. As expected, participants in the white-ideograph condition did not judge the white ideographs to translate positive words correctly more often ($M = 2.84$, $SD = 1.74$) than for negative words ($M = 2.53$, $SD = 1.39$), $t(18) = 0.62$, $p = .54$, whereas participants in the black-ideograph condition inferred that the black ideographs were the correct translation for positive words ($M = 2.35$, $SD = 0.99$) less often than for negative words ($M = 3.80$, $SD = 1.24$), $t(19) = 4.31$, $p < .001$.[4]

---

[4] Additional simple comparisons across the between-participant factor (brightness) similarly show that participants translated negative words more often with black than white ideographs, $t(37) = 3.02$, $p = .005$, whereas the choices to translate positive words by white or black ideographs did not differ significantly, $t(37) = 1.09$, $p = .281$.

Discussion

When brightness was manipulated between participants, white ideographs did not bias translation judgments for positive or negative words, whereas black ideographs did influence translation judgments for positive and negative words. While the association between black and negativity was clearly hypothesized, the non-association between black and negativity (positive is not black) is not a necessity but follows from the assumption that black is associated with negativity. These findings are in line with the importance of the activation of shared relational structures in the brightness and valence domains, which predicts that the positivity of white emerges only in opposition to the negativity of black. The differences between Experiment 3 and Experiment 1 are not a priori predicted by a one-to-one grounding account.

Experiment 4

We proceeded with a final comparison of the predictions made by the shared relational structures view and the one-to-one grounding account by manipulating both the valence and brightness dimensions between participants. In addition, two baseline conditions were added where participants were asked to translate neutral words by either white or black ideographs. These conditions were not expected to reveal a deviation from guessing average in the translation judgments, to provide support for our assumption that absent any brightness–valence associations, *yes* and *no* translation judgments are equally likely to occur. Given the stable association between black and negativity, only translation judgments for negative words by black ideographs were expected to show a deviation from guessing average, with no expected differences for the white-ideograph positive-word condition (or any of the other conditions) due to the absence of shared relational structures.

Method

Participants and procedure. A total of 122 students (74 females, mean age 20 years) at a Dutch university took part in this experiment for monetary compensation (€3, or $4.50) and were randomly assigned to one of the six conditions of the 3 (word valence: positive vs. negative vs. neutral) $\times$ 2 (ideograph brightness: white vs. black) between-participants design. The procedure was identical to Experiment 3, with the exception that the valence of the 12 stimulus words was manipulated between participants.

Results

As expected based on the negative evaluations of the color black, the average number of *yes* judgments was significantly

higher than guessing average in the negative-word black-ideograph condition but did not differ significantly in the other five conditions (for averages and *t* tests, see Table 1). Although the interaction between word valence and ideograph color did not reach significance ($p > .10$), theory-driven planned contrast confirmed our prediction that *yes* judgments to translate negative words by black ideographs were not only above guessing average but also higher than the average *yes* judgments for the other conditions, $t(116) = 1.93$, $p = .028$, one-sided.

Discussion

The results of Experiment 4, together with the findings from Experiments 1–3, confirm that white ideographs are not judged to correctly translate the meaning of positive words above guessing average unless the perceptual or conceptual light–dark and positive–negative oppositions are present. On the other hand, black ideographs are always judged to correctly translate the meaning of negative words above guessing average (see Figure 2, top panel). This latter result is also important from a methodological perspective, since it shows that the Chinese translation paradigm can be used to reveal the influence of perceptual characteristics of the ideographs on the likelihood with which they are seen to correctly translate valenced words, even without activating opposite endpoints of the conceptual or perceptual dimensions. Finally, the current experiment revealed no default bias in translation judgments for neutral stimulus words, supporting our assumption that statistically testing for deviations from guessing average is a valid approach.

Experiment 5

The first four experiments reveal the importance of shared relational structures for the association between positivity and white to emerge. When manipulating the brightness or valence dimensions between participants, positive words are no longer translated above guessing average by white ideographs. Obviously, when stimulus dimensions are manipulated between individuals, the positive–negative and/or white–black opposition is no longer salient. The literature suggests that another way to reduce the salience of oppositions in the stimulus dimension in many experimental paradigms is to add neutral stimuli (e.g., Duscherer et al., 2008; Everaert, Spruyt, & De Houwer, 2011; Kahneman & Chajczyk, 1983). If the mapping effects observed in Experiment 1 occurred due to the presence of a binary opposition in the valence and brightness dimension, adding a third category of neutral stimuli to the task should dilute the salience of shared relational

Table 1

Average Number of Yes Judgments and One-Sample t *Tests Against Guessing Average for Each Condition in Study 4*

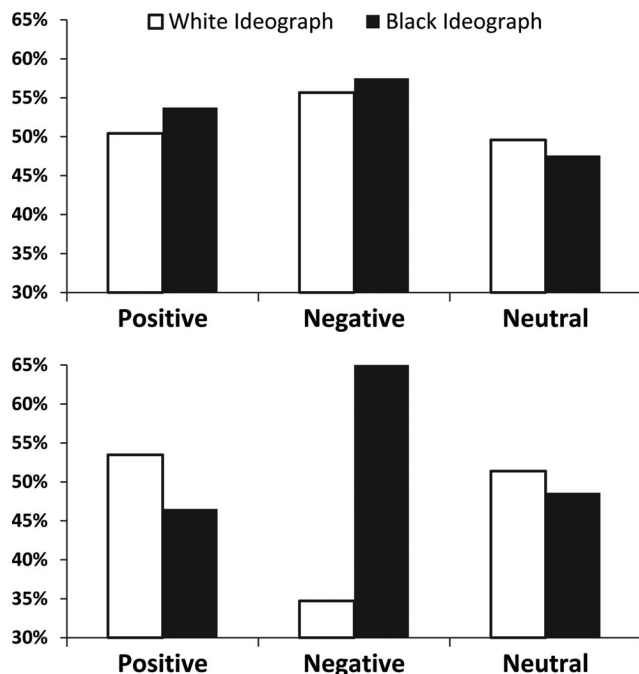| Color	Word valence	M	%	SD	t	df	p	d
White	Positive	6.05	50.42	1.50	0.15	20	.89	.03
White	Negative	6.68	55.67	1.70	1.75	18	.10	.40
White	Neutral	5.95	49.58	1.40	0.87	19	.87	.04
Black	Positive	6.45	53.75	1.91	1.06	19	.30	.24
Black	Negative	6.95	57.92	1.15	3.71	19	.00	.80
Black	Neutral	5.71	47.58	1.79	0.73	20	.47	.16

Figure 2. Percentages of *yes* judgments (Experiment 4, top panel) or black- versus white-ideograph choices (mutually exclusive in Experiment 5, bottom panel) compared to guessing average (50%).

structures, reducing the likelihood that the positivity of white emerges.

Method

Participants and procedure. Twenty-four students (15 males, mean age 22 years) at a Dutch university took part in this experiment for monetary compensation (€3, or $4.50). The procedure was identical to Experiment 1, with the exception that six neutral words were added to the translation task, which now consisted of 18 translation judgments.

Results

As in Experiment 1, participants translated the six negative words below guessing average by white ideographs (34.72%, $M = 2.08$, $SD = 1.14$), as indicated by a test against chance, $t(23) = 3.94$, $p = .001$, Cohen's $d = .81$. Unlike Experiment 1, the six positive words were not translated above guessing average by white ideographs (53.47%, $M = 3.21$, $SD = 1.06$): Although the average translation judgments were descriptively in the same direction as in Experiment 1, the translation bias did not reach statistical significance, $t(23) = 0.96$, $p = .35$, $d = .20$. As in Experiment 4, no translation bias was observed for neutral stimuli (51.39%, $M = 3.08$, $SD = 1.44$), $t(23) = 0.28$, $p = .78$, $d = .06$. A repeated measures ANOVA with word valence as the only within-subject factor indicated that the translation choices for the white ideograph significantly differed depending upon the valence of the stimulus words, $F(2, 46) = 5.09$, $p = .010$, $\eta_p^2 = .18$. Paired-sample t tests revealed that participants chose the white ideograph less to translate negative words compared to positive,

$t(23) = 3.24$, $p = .004$, and neutral words, $t(23) = 2.26$, $p = .034$, which did not differ from each other, $t(23) = 0.35$, $p = .73$ (see Figure 2, bottom panel).

Discussion

After replacing the binary opposition in the valence dimension (positive vs. negative) by three ordinally related categories (positive, neutral, and negative), no mapping between positivity and white emerged during the translation task. This finding provides further support for the importance of structural relations consisting of bipolar opposites between the valence and brightness dimensions. The current finding suggests that in addition to the within versus between manipulation of the dimensions under investigation (Experiments 1–4), another important structural factor in experimental designs is whether the stimulus dimensions are both manipulated dichotomously or not. The current experiment is the first to suggest that reducing the salience of the binary oppositions in the stimulus dimension by adding neutral stimuli to the valence opposition is enough to remove structural mapping effects in experimental paradigms.

Experiment 6

After conceptually replicating earlier findings by Meier et al. (2004) in Experiment 1, the subsequent Experiments 2–5 supported the shared relational structures view by showing no deviations from guessing average for the average number of translations of positive words by white ideographs. In the current experiment, we aimed to activate the positivity of white not by manipulating both endpoints of the valence and brightness dimensions within participants as in Experiment 1 but by presenting positive words and white ideographs in perceptual opposition to the negativity of black. Black and negativity were activated by a compound cue consisting of a black background (for a conceptually related use of backgrounds as a contextually salient cue, see Gawronski, Rydell, Vervliet, & De Houwer, 2010). As the pilot studies (see the online supplemental materials) have shown, a black background is evaluated negatively. Presenting positive words and white ideographs on a black background should therefore activate the shared opposition in the valence and brightness dimensions. The current experiment tested the prediction that translation judgments for positive words by white ideographs presented on a gray background would not differ from guessing average (replicating Experiment 4). However, when white ideographs are presented on a black background, we hypothesized that participants would judge the white ideographs as the correct translation for positive words above guessing average. This prediction follows only from a shared relational structures view.

Method

Participants and procedure. Forty-eight university students (32 females, mean age 21 years) took part in this experiment for monetary compensation (€3, or $4.50) and were randomly assigned to either the black or gray background condition. The procedure was identical to the positive-words white-ideographs condition in Experiment 4. However, whereas, for half the participants, the white ideographs and positive words were presented

against a gray background (identical to that of Experiments 1–5), the white ideographs were presented against a black background for the other half of the participants.

Results

The average number of times participants answered *yes* to the question whether the white ideograph correctly translated the positive words was calculated. Whereas this average (49.33%, $M = 5.92$, $SD = 1.19$) did not differ from guessing average in the gray background condition, $t(24) = -0.34$, $p = .74$, Cohen's $d = .07$, the average number of *yes* judgments was higher than guessing average when the white ideograph was presented against a black background (54.33%, $M = 6.52$, $SD = 0.90$), $t(22) = 2.79$, $p = .01$, Cohen's $d = .58$. In line with our prediction, the average number of *yes* translation was higher in the black background condition compared to the gray background condition, $t(46) = 1.97$, $p = .028$, one-tailed.

Discussion

Presenting white ideographs and positive words against a black background increases the likelihood that participants will judge the ideograph to correctly translate the positive word. This effect is not present when the translation task is performed against a gray background. As predicted from the shared relational structures view, the positivity of white emerges after introducing the opposition of the negativity of the black background compared to the positive words and white ideographs. These findings show how subtle task characteristics such as the color of the background can determine whether shared relational structures between stimulus dimensions will emerge or not.

General Discussion

The current studies examined whether brightness is intrinsically positive and darkness is intrinsically negative by comparing a one-to-one grounding account for the valence–brightness mapping with a shared relational structures view. Over six experiments, our results revealed that whereas black is associated with negativity independent of contextual factors (see also Scherer & Lambert, 2011), the positivity of white emerges only in direct opposition with the negativity of black (see Figures 1 and 2). As such, the affordance of white to represent positivity is dependent upon the within-participant manipulation of both endpoints of the brightness and valence dimensions (Experiment 1) or upon the perceptual activation of the light–dark and positive–negative opposition (Experiment 6). On the basis of the present results, we conclude that a one-to-one grounding account cannot fully explain the use of white as a representation of positivity and that the activation of the shared relational structures in the valence and brightness domains is essential for the positivity–white mapping to emerge. Importantly, the mapping between white and positivity is the result of structural factors of the experimental design (e.g., the within-participant manipulation of stimulus dimensions) but does not necessarily reflect stable associations represented in memory.

The importance of activating the polar oppositions in the brightness and valence dimensions is not at odds with the assumption that the brightness–valence relationship is automatic (Meier et al.,

2004; Sherman & Clore, 2009). The all-or-none conception of automaticity has been challenged in recent years (for a review, see Moors & De Houwer, 2006), and many processes that were previously assumed to be fully automatic, such as the Stroop effect argued to underlie metaphor congruency effects (Meier & Robinson, 2004; Meier et al., 2004; Schubert, 2005; Sherman & Clore, 2009), have turned out to be more conditional and less automatic than assumed (e.g., Kahneman & Chajczyk, 1983; Spruyt, De Houwer, Hermans, & Eelen, 2009). The association between white and positivity might be better regarded as conditionally automatic, depending on the activation of the shared opposition between light–dark and good–bad. Given that the polar opposites in these dimensions are salient enough (e.g., by manipulating the endpoints of these dimensions within subjects), the brightness–valence association could influence judgments automatically (Meier et al., 2004; Sherman & Clore, 2009).

Interestingly, the negativity of black was strong enough to influence translations judgments in the first five experiments irrespective of whether polar oppositions were present or absent. Scherer and Lambert (2011) similarly found that evaluations for Black faces were negative independent of the context but that the positivity of White faces was context dependent. A possible explanation for the observation that the positive meaning of white only emerges in opposition to the negativity of black, whereas black always activates negativity, might be found in the fact that the meaning of colors can be grounded through evolutionary predispositions on the one hand and culturally learned associations through repeated pairings of a color with a specific concept on the other hand (Elliot & Maier, 2007). The negativity of black shows little cultural variation (Adams & Osgood, 1973) and has its theorized origin in the fact that humans are diurnal animals (Grillon et al., 1997; Williams & Morland, 1976). The valence of white, on the other hand, seems to reflect culturally learned associations (Saito, 1996) instead of biological predispositions. White can represent positivity, but this affective meaning does not seem to emerge when bright stimuli are encountered in isolation. However, if the relational structure of the shared oppositions in the valence and brightness dimensions are activated, then the positive–bright mapping can bias judgments to the same extent as the negative–dark mapping (Experiment 1).

Although speculative, darkness might have more often led directly to negative affect in daily life, while positive experiences in light situations are less often directly related to the brightness. Brightness as a perceptual representation of positivity might be a culturally learned metaphor, perhaps often used in opposition to the negativity of black, such as when the white knight in shining armor battles the black knight. Future research could explore the circumstances under which the activation of shared relational structures is necessary for metaphoric mappings to emerge or when concrete experiences and abstract concepts are associated irrespective of the context or their relational structures (see also Schneider, Rutjens, Jostmann, & Lakens, 2011). Note that the presence of shared relational structures is only one way in which the positivity of white can emerge. Other contextual factors such as specific task instructions or activating a certain mind-set may similarly allow the positivity of white to emerge.

In addition to providing new insights into the processes underlying the metaphoric representation of valence in brightness, the current set of studies has important methodological implications

for paradigms that aim to understand the intrinsic meaning of stimuli by investigating automatic associations. As our studies reveal, biases in participants' responses do not necessarily reflect stable associations but can be the result of structural factors in experimental designs, such as shared relational structures. Conceptually related investigations have recently revealed several structural factors that underlie biases in judgments under uncertainty as well as reaction times in speeded categorization tasks, such as contrast effects (Scherer & Lambert, 2009, 2011), polarity differences (Lakens, 2011; Proctor & Cho, 2006), evaluative response coding (Eder & Rothermund, 2008), and salience asymmetries (Rothermund & Wentura, 2004). What all these findings have in common is that observed relationships between stimuli (or stimuli and responses) are not caused by intrinsic associations but by structural relationships between stimuli in the experimental task.

Shared relational structures might play a role in other experimental paradigms that rely on choices between a white and black choice alternative. For example, consider the Color Meaning Test II (CMT II; Williams, Boswell, & Best, 1975) and the Preschool Racial Attitude Measure II (PRAM II; Williams, Best, Boswell, Mattson, & Graves, 1975) that have often been used to measure a pro-white/anti-black bias in children. Children read a sentence with a positive or negative adjective and are asked to choose between a white or black animal or human (e.g., Which is the bad man?). Although the preference to choose white figures for positive behaviors and black figures for negative behaviors is treated and interpreted as a unified racial bias, the results presented here show that these studies do not necessarily tap into the intrinsic meaning assigned to the colors black and white, given that children, just like adults, rate the color white as neutral (Burkitt et al., 2003; Meerum Terwogt & Hoeksma, 1995). We do not mean to cast doubt on the presence of racial stereotypes, which has been observed using many different paradigms (e.g., Degner & Wentura, 2010; Smith-McLallen et al., 2006), but want to highlight the possibility that structural factors in experimental designs influenced the results observed in these studies. A recent study, conceptually related to the current investigation, revealed that structural factors unrelated to implicit attitudes do indeed play a role when investigating racial stereotypes. Scherer and Lambert (2011) found that attitudes toward Black faces were consistently negative irrespective of the context. Attitudes toward White faces were positively biased only when presented in the same block as strongly negative target stimuli (i.e., pictures of attacking animals).

The current observations add to an increasing awareness that cognition is inherently contextual. As eloquently formulated by Bruner (1990, p. 64),

> we shall be able to interpret meanings and meaning-making in a principled manner only in the degree to which we are able to specify the structure and coherence of the larger contexts in which specific meanings are created and transmitted.

Recent theoretical developments in psychology have rekindled an interest in more contextual views on psychological meaning (Mesquita, Barrett, & Smith, 2010), and researchers are providing new ways to think about how meaning emerges (see Barsalou, 2008; Clark, 2008; Smith & Semin, 2004). The current findings highlight the inherently contextualized construction of meaning structures.

As a consequence, associations should be interpreted in the relational context in which they are observed.

References

Adams, F. M., & Osgood, C. E. (1973). A cross-cultural study of the affective meanings of color. *Journal of Cross-Cultural Psychology, 4,* 135–156. doi:10.1177/002202217300400201

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology, 59,* 617–645. doi:10.1146/annurev.psych.59.103006.093639

Boroditsky, L. (2000). Structural similarity: Understanding time through spatial metaphors. *Cognition, 75,* 1–28. doi:10.1016/S0010-0277(99)00073-6

Bruner, J. (1990). *Acts of meaning.* Cambridge, MA: Harvard University Press

Burkitt, E., Barrett, M., & Davis, A. (2003). Children's colour choices for completing drawings of affectively characterised topics. *Journal of Child Psychology and Psychiatry, 44,* 445–455. doi:10.1111/1469-7610.00134

Cacioppo, J. T., Priester, J. R., & Berntson, G. G. (1993). Rudimentary determinants of attitudes: II. Arm flexion and extension have differential effects on attitudes. *Journal of Personality and Social Psychology, 65,* 5–17. doi:10.1037/0022-3514.65.1.5

Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension.* Oxford, England: Oxford University Press.

Crawford, L. E. (2009). Conceptual metaphors of affect. *Emotion Review, 1,* 129–139. doi:10.1177/1754073908100438

Degner, J., & Wentura, D. (2010). Automatic prejudice in childhood and early adolescence. *Journal of Personality and Social Psychology, 98,* 356–374. doi:10.1037/a0017993

Deutsch, R., & Gawronski, B. (2009). When the method makes a difference: Antagonistic effects on "automatic evaluations" as a function of task characteristics of the measure. *Journal of Experimental Social Psychology, 45,* 101–114. doi:10.1016/j.jesp.2008.09.001

Duscherer, K., Holender, D., & Molenaar, E. (2008). Revisiting the affective Simon effect. *Cognition & Emotion, 22,* 193–217. doi:10.1080/02699930701339228

Eder, A. B., & Rothermund, K. (2008). When do motor behaviors (mis)match affective stimuli? An evaluative coding view of approach and avoidance reactions. *Journal of Experimental Psychology: General, 137,* 262–281. doi:10.1037/0096-3445.137.2.262

Elliot, A. J., & Maier, M. A. (2007). Color and psychological functioning. *Current Directions in Psychological Science, 16,* 250–254. doi:10.1111/j.1467-8721.2007.00514.x

Everaert, T., Spruyt, A., & De Houwer, J. (2011). On the (un)conditionality of automatic attitude activation: The valence proportion effect. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale, 65,* 125–132. doi:10.1037/a0022316

Frank, M. G., & Gilovich, T. (1988). The dark side of self- and social perception: Black uniforms and aggression in professional sports. *Journal of Personality and Social Psychology, 54,* 74–85. doi:10.1037/0022-3514.54.1.74

Gattis, M. (2001). *Spatial schemas and abstract thought.* Cambridge, MA: MIT Press.

Gawronski, B., Cunningham, W. A., LeBel, E. P., & Deutsch, R. (2010). Attentional influences on affective priming: Does categorization influence spontaneous evaluations of multiply categorizable objects? *Cognition & Emotion, 24,* 1008–1025. doi:10.1080/02699930903112712

Gawronski, B., Rydell, R. J., Vervliet, B., & De Houwer, J. (2010). Generalization versus contextualization in automatic evaluation. *Journal of Experimental Psychology: General, 139,* 682–701. doi:10.1037/a0020315

Gentner, D., & Bowdle, B. (2008). Metaphor as structure mapping. In R. W. Gibbs (Ed.), *Cambridge handbook of metaphor and thought* (pp. 109–128). Cambridge, England: Cambridge University Press.

Götz, K. O., & Götz, K. (1974). Color attitudes of art students and university students: I. Imagined colors. *Perceptual and Motor Skills, 38,* 63–70. doi:10.2466/pms.1974.38.1.63

Grillon, C., Pellowski, M., Merikangas, K. R., & Davis, M. (1997). Darkness facilitates acoustic startle reflex in humans. *Biological Psychiatry, 42,* 453–460. doi:10.1016/S0006-3223(96)00466-0

Kahneman, D., & Chajczyk, D. (1983). Test of the automaticity of reading: Dilution of Stroop effects by color-irrelevant stimuli. *Journal of Experimental Psychology: Human Perception and Performance, 9,* 497–509. doi:10.1037/0096-1523.9.4.497

Lakens, D. (2011). Polarity correspondence in metaphor congruency effects: Structural overlap predicts categorization times for bipolar concepts presented in vertical space. *Journal of Experimental Psychology: Learning, Memory, and Cognition.* Advance online publication. doi:10.1037/a0024955

Lakens, D., Semin, G. R., & Foroni, F. (2011). Why your highness needs the people: Comparing the absolute and relative representation of power in vertical space. *Social Psychology, 42,* 205–213. doi:10.1027/1864-9335/a000064

Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (Ed.), *Metaphor and thought* (pp. 202–251). Cambridge, England: Cambridge University Press.

Lakoff, G. (2008). The neural theory of metaphor. In R. W. Gibbs (Ed.), *Cambridge handbook of metaphor and thought* (pp. 17–38). Cambridge, England: Cambridge University Press.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by.* Chicago, IL: University of Chicago Press.

Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to western thought.* Chicago, IL: University of Chicago Press.

McManus, I. C., Jones, A. L., & Cottrell, J. (1981). The aesthetics of colour. *Perception, 10,* 651–666. doi:10.1068/p100651

Meerum Terwogt, M., & Hoeksma, J. B. (1995). Colors and emotions: Preferences and combinations. *Journal of General Psychology, 122,* 5–17. doi:10.1080/00221309.1995.9921217

Meier, B. P., & Robinson, M. D. (2004). Why the sunny side is up: Associations between affect and vertical position. *Psychological Science, 15,* 243–247. doi:10.1111/j.0956-7976.2004.00659.x

Meier, B. P., Robinson, M. D., & Clore, G. L. (2004). Why good guys wear white: Automatic inferences about stimulus valence based on brightness. *Psychological Science, 15,* 82–87. doi:10.1111/j.0963-7214.2004.01502002.x

Mesquita, B., Barrett, L. F., & Smith, E. R. (Eds.). (2010). *The mind in context.* New York, NY: Guilford Press.

Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin, 132,* 297–326. doi:10.1037/0033-2909.132.2.297

Moors, A., Spruyt, A., & De Houwer, J. (2010). In search of a measure that qualifies as implicit. In K. Payne & B. Gawronski (Eds.), *Handbook of implicit social cognition: Measurement, theory, and applications* (pp. 19–37). New York, NY: Guilford Press.

Murphy, G. (1996). On metaphoric representation. *Cognition, 60,* 173–204.

Murphy, S. T., & Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology, 64,* 723–739. doi:10.1037/0022-3514.64.5.723

Osgood, C. E. (1952). The nature and measurement of meaning. *Psychological Bulletin, 49,* 197–237. doi:10.1037/h0055737

Osgood, C. E. (1960). The cross-cultural generality of visual-verbal synesthetic tendencies. *Behavioral Science, 5,* 146–169. doi:10.1002/bs.3830050204

Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The measurement of meaning.* Urbana: University of Illinois Press.

Payne, B. K., Cheng, S. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology, 89,* 277–293.

Proctor, R. W., & Cho, Y. S. (2006). Polarity correspondence: A general principle for performance of speeded binary classification tasks. *Psychological Bulletin, 132,* 416–442. doi:10.1037/0033-2909.132.3.416

Rothermund, K., & Wentura, D. (2004). Underlying processes in the implicit association test: Dissociating salience from associations. *Journal of Experimental Psychology: General, 133,* 139–165. doi:10.1037/0096-3445.133.2.139

Saito, M. (1996). A comparative study of color preferences in Japan, China and Indonesia, with emphasis on the preference for white. *Perceptual and Motor Skills, 83,* 115–128. doi:10.2466/pms.1996.83.1.115

Scherer, L. D., & Lambert, A. J. (2009). Contrast effects in priming paradigms: Implications for theory and research on implicit attitudes. *Journal of Personality and Social Psychology, 97,* 383–403. doi:10.1037/a0015844

Scherer, L. D., & Lambert, A. J. (2011). Implicit race bias revisited: On the utility of task context in assessing implicit attitude strength. *Journal of Experimental Social Psychology.* Advance online publication. doi:10.1016/j.jesp.2011.06.010

Schneider, I. K., Rutjens, B., Jostmann, N. B., & Lakens, D. (2011). Weighty matters: Importance literally feels heavy. *Social Psychological & Personality Science, 2,* 474–478. doi:10.1177/1948550610397895

Schubert, T. W. (2005). Your highness: Vertical positions as perceptual symbols of power. *Journal of Personality and Social Psychology, 89,* 1–21. doi:10.1037/0022-3514.89.1.1

Schwarz, N. (2007). Attitude construction: Evaluation in context. *Social Cognition, 25,* 638–656.

Sherman, G. D., & Clore, G. L. (2009). The color of sin: White and black are perceptual symbols of moral purity and pollution. *Psychological Science, 20,* 1019–1025. doi:10.1111/j.1467-9280.2009.02403.x

Smith, E. R., & Semin, G. R. (2004). Socially situated cognition: Cognition in its social context. *Advances in Experimental Social Psychology, 36,* 53–117. doi:10.1016/S0065-2601(04)36002-8

Smith-McLallen, A., Johnson, B. T., Dovidio, J. F., & Pearson, A. R. (2006). Black and white: The role of color bias in implicit race bias. *Social Cognition, 24,* 46–73. doi:10.1521/soco.2006.24.1.46

Spruyt, A., De Houwer, J., Hermans, D., & Eelen, P. (2007). Affective priming of nonaffective semantic categorization responses. *Experimental Psychology, 54,* 44–53. doi:10.1027/1618-3169.54.1.44

Williams, J. E., Best, D. L., Boswell, D. A., Mattson, L. A., & Graves, D. J. (1975). Preschool Racial Attitude Measure II. *Educational and Psychological Measurement, 35,* 3–18. doi:10.1177/001316447503500101

Williams, J. E., Boswell, D. A., & Best, D. L. (1975). Evaluative responses of preschool children to the colors white and black. *Child Development, 46,* 501–508. doi:10.2307/1128148

Williams, J. E., & Morland, J. K. (1976). *Race, color, and the young child.* Chapel Hill: University of North Carolina Press.